October 2001

Science & Technology Review

NNSA
U.S. Department of Energy's
Lawrence Livermore
National Laboratory

Supercomputing Becomes an Institutional Resource

Also in this issue:

- More Advances in Ultrashort-Pulse Lasers
- Modeling Dispersions of Biological and Chemical Agents
- Centennial of E. O. Lawrence's Birth

About the Cover

Computing systems leader Greg Tomaschke works at the console of the 680-gigaops Compaq TeraCluster2000 parallel supercomputer, one of the principal machines used to address large-scale scientific simulations at Livermore. The supercomputer is accessible to unclassified program researchers throughout the Laboratory, thanks to the Multiprogrammatic and Institutional Computing (M&IC) Initiative described in the article beginning on p. 4. M&IC makes supercomputers an institutional resource and helps scientists realize the potential of advanced, three-dimensional simulations.

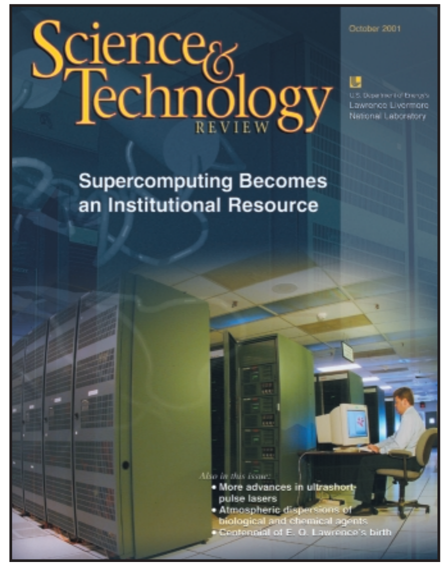

Cover design: Amy Henke

About the Review

Lawrence Livermore National Laboratory is operated by the University of California for the Department of Energy's National Nuclear Security Administration. At Livermore, we focus science and technology on assuring our nation's security. We also apply that expertise to solve other important national problems in energy, bioscience, and the environment. *Science & Technology Review* is published 10 times a year to communicate, to a broad audience, the Laboratory's scientific and technological accomplishments in fulfilling its primary missions. The publication's goal is to help readers understand these accomplishments and appreciate their value to the individual citizen, the nation, and the world.

Please address any correspondence (including name and address changes) to *S&TR*, Mail Stop L-664, Lawrence Livermore National Laboratory, P.O. Box 808, Livermore, California 94551, or telephone (925) 423-3432. Our e-mail address is str-mail@llnl.gov. *S&TR* is available on the World Wide Web at www.llnl.gov/str/.

National Nuclear Security Administration

Prepared by **LLNL** under contract No. W-7405-Eng-48

Science & Technology REVIEW

October 2001

Lawrence
Livermore
National
Laboratory

Contents

S&TR Staff

Scientific Editor
Andrew A. Quong

Managing Editor
Ray Marazzi

Publication Editor
Gloria Wilt

Writers
Arnie Heller, Ann Parker,
Katie Walter, and Gloria Wilt

Art Directors
Amy Henke and George Kitrinos

Designer
Amy Henke

Internet Designer
Kitty Tinsley

Compositor
Louisa Cardoza

Proofreader
Carolin Middleton

S&TR, a Director's Office publication, is produced by the Technical Information Department under the direction of the Office of Policy, Planning, and Special Studies.

S&TR is available on the Web at www.llnl.gov/str/.

Printed in the United States of America

Available from
National Technical Information Service
U.S. Department of Commerce
5285 Port Royal Road
Springfield, Virginia 22161

UCRL-52000-01-10
Distribution Category UC-0
October 2001

Features

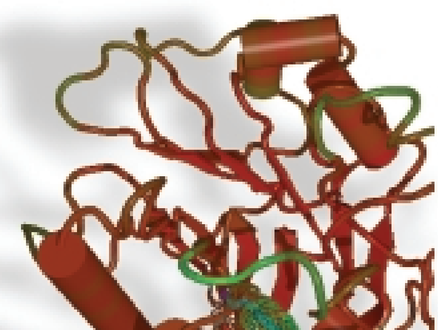

Research Highlights

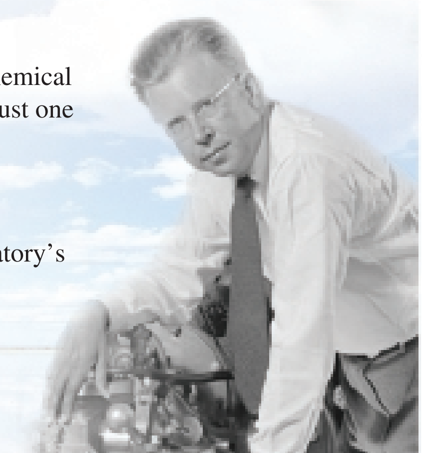

Departments

Lab represents U.S. in nuclear waste study

Five countries and regions have agreed in principle to participate in a joint research project on deep underground disposal of spent nuclear fuel. Work to develop disposal technologies would be performed at the research centers of the five participants: Lawrence Livermore in the U.S., the Japan Nuclear Cycle Development Institute, the (South) Korean Atomic Energy Research Institute, the Beijing Research Institute of Uranium Geology in the People's Republic of China, and the Institute of Nuclear Energy Research in Taiwan.

The U.S. Department of Energy is playing a leading role in the project and expects to build a joint research center near Las Vegas, Nevada. Says C. K. Chou, associate director of the Energy and Environment Directorate at Livermore, "The United States has already spent about $5 billion for an underground disposal project at Yucca Mountain, Nevada. We want to take advantage of the knowledge gained. Nations with advanced nuclear power generation technology have a responsibility to propose a way to dispose of spent nuclear fuel, while also promoting nuclear reactor safety."

The joint research project must solve a number of technical problems, such as what type of rock is most suitable for a nuclear waste repository and how to prevent radioactive material from contaminating groundwater. South Korean officials have proposed that research be conducted at an underground experimental facility they plan to build in Seoul. Officials from the Japan Nuclear Cycle Development Institute and Lawrence Livermore have proposed contributing academic papers and computer programs.

Although there is a possibility that China may accept spent nuclear fuel from Taiwan, the project is not expected to deal with the disposal policies of the participating nations, some of whom dispose of spent fuel within their borders while others send it to other countries.

Contact: C. K. Chou (925) 422-4950 (chou1@llnl.gov).

Predicting how wildfires will burn

In the fire season of 2000, wildfires burned 6.8 million acres of public and private lands, including large parts of Los Alamos National Laboratory. Experts believe that annual wildfires will increase and will ravage thousands of acres of land and endanger human life.

In response to the threat, scientists at the Lawrence Livermore and Los Alamos national laboratories are working on an initiative for the National Wildfire Prediction Program. They are combining Los Alamos's multiyear wildfire modeling effort with existing capabilities at the National Atmospheric Release Advisory Center at Livermore to predict the behavior of wildfires and prescribed burns. The intention is to provide around-the-clock guidance to fire management planners for the most effective use of firefighting resources. They also want to predict the behavior of fires of strategic interest around the globe.

The Livermore–Los Alamos team has already developed wildfire models and accurately simulated the behavior of historic fires. Livermore researchers are linking a fire model to a regional weather prediction model and performing simulations to reconstruct the early stage of the 1991 fire in the Oakland hills of northern California. The simulations test the combined model's accuracy. Follow-on studies, requested by emergency management and planning officials, will look at hypothetical fires in nearby canyons that escaped the 1991 fire. The information will improve preparedness for future wildfires.

Livermore atmospheric scientist Michael Bradley says that results of the modeling and prediction initiative could move the nation into a new era of scientifically based wildfire and vegetation management. He sees the day when fire trucks will carry laptops to tap into a national wildfire behavior prediction center and determine where to direct firefighting troops. Eventually, the models might even predict the effects of firefighting activities, which means firefighters would be able to choose the safest and most effective techniques for specific fires.

Contact: Michael Bradley (925) 422-1835 (bradley6@llnl.gov).

Site 300 gets new lightning warning system

Lightning strikes there are rare, but because Site 300 functions as an explosives test facility, conservative safety precautions for lightning events are warranted. Thus, the site has upgraded its lightning detection and warning system. The system uses electric field mils, which detect the strength of the electric field gradient, thereby indicating the potential for lightning, and a new electrical storm identification device, which optically detects and measures the site's distance from lightning flashes.

Larry Sedlacek, Site 300 manager, says, "The new detection equipment improves our ability to accurately detect potential lightning conditions and safely shut down explosives operations during those times." He adds, "We'll be better able to gauge when employees working in the field need to evacuate to a protected building."

Safety procedures during lightning events have been upgraded as well. The procedures describe conditions that determine the Site 300 lightning status, which is designated as all clear, lightning watch, or lightning alert. Each designation is associated with appropriate actions. Employees at Site 300 are warned of lightning alerts by building, radio, and alpha pages. The complete procedures and more lightning information can be found on the Web at www.llnl.gov/site300/PDFs/Final_Lightning.pdf.

Contact: Larry Sedlacek (925) 422-8853 (sedlacek2@llnl.gov).

Commentary by Dona Crawford

Supercomputing Resources Are Vital to Advancing Science

FOR several decades, the Computation Directorate has provided leading-edge computational capability to support the Laboratory's missions and programs. Increasingly, these programs rely on computer simulations to guide scientific discovery and engineering development. Our goal in Computation is to provide and, where necessary, develop the computational and information technologies that enable Laboratory programs to fully realize the promise of scientific simulations, especially those generated with three-dimensional models.

More and more, advanced simulations require the power of parallel supercomputers that employ hundreds or even thousands of processors working in tandem. The most advanced parallel supercomputers are those developed for the Accelerated Strategic Computing Initiative (ASCI), an important element of the nation's Stockpile Stewardship Program to assure the safety and reliability of the U.S. nuclear stockpile. Last year, Livermore installed the world's most powerful supercomputer, the 12-teraops (trillion operations per second) ASCI White system manufactured by IBM.

In the mid-1990s, Livermore launched a bold initiative to make available to all unclassified researchers a computing environment similar to that pioneered by ASCI. This effort, described in the article beginning on p. 4, is called the Multiprogrammatic and Institutional Computing (M&IC) Initiative and represents a strong partnership between the Laboratory and its research programs. We launched this initiative because desktop computers, while remarkably capable compared to their predecessors of just a few years ago, do not provide the computational power of centralized systems, which is needed for simulations of the highest resolution.

As the article details, the M&IC Initiative has sponsored the acquisition of increasingly powerful platforms, such as the TeraCluster2000 supercomputer manufactured by Compaq Computer Corporation. This acquisition benefited from a Cooperative Research and Development Agreement between Compaq and the Laboratory. M&IC has also made it possible to access limited, unclassified portions of our ASCI supercomputing platforms made by IBM. Other resources include a number of smaller machines for numerous smaller simulation runs.

Hardware is only part of the story. We've made a strong effort to provide tools and techniques for analyzing and improving the performance of simulation codes on massively parallel computers and new ways to visualize the large amounts of data that are routinely generated. The capabilities allow researchers to plan and efficiently run their simulations. We've also forged productive partnerships with leading equipment makers such as IBM and Compaq, and we maintain strong interactions with the academic scientific computing community to stay abreast of this rapidly evolving field.

The new computational resources are enabling Livermore scientists to conduct unclassified simulations on a scale that used to be considered impossible. The simulations include, to name a few, exploring the interiors of stars, analyzing the response of materials to extreme pressures and temperatures, studying the consequences of global warming and climate change, understanding the mechanisms of genes and proteins, and investigating the effects of seismic forces on structures. Such simulations, done at unprecedented resolution, are leading to important new discoveries and insights.

As a result of increasing demand, we are considering new options to strengthen our unclassified computing resources and make them even more cost-effective. We're confident that Livermore will continue its long history of preeminence in providing leading-edge computational capability to its missions and programs.

■ Dona Crawford is associate director for Computation.

Sharing the Power of Supercomputers

More researchers are getting ready access to Livermore's powerful supercomputers for advanced simulations.

COMPUTER simulation is changing the nature of scientific discovery by becoming a full partner with theory and experiment. Nowhere is that transformation more visible than at Lawrence Livermore, where scientists are relying increasingly on computer simulations, especially those based on high-resolution, three-dimensional models. Such simulations, however, require the enormous computational horsepower of the latest generation of supercomputers.

To ensure that all Livermore programs and researchers have the possibility of accessing high-performance computers, the Laboratory created the Multiprogrammatic and Institutional Computing (M&IC) Initiative in 1996. The initiative has grown substantially in the last few years and currently serves more than a thousand users, including outside collaborators, with some of the most powerful computers available. Indeed, virtually every Livermore unclassified program—from atmospheric sciences to bioscience—benefits from the computing resources provided by the initiative.

"M&IC is a partnership between the Laboratory and its research programs to bring high-performance computing to every researcher," says Mike McCoy, head of Livermore's Scientific Computing and Communications Department. M&IC recognizes that no matter what their mission, scientists should have ample access to centralized, high-performance computers far more capable than the computational resources that individual departments could afford to purchase and support. M&IC leaders, under the direction of McCoy and former Computation Associate Director Dave Cooper, have worked with key researchers around the Laboratory to create a centralized operation that helps all researchers do better science through more advanced simulations.

McCoy notes that his department's goal is to make the "discovery environment" that is evolving in the Accelerated Strategic Computing Initiative (ASCI) available to all researchers. ASCI is an element of the National Nuclear Security Administration's Stockpile Stewardship Program to assure the safety and reliability of the nation's nuclear weapons in the absence of underground testing. (See *S&TR*, June 2000, pp. 4–14.) In the ASCI computing approach, scientists are supported by powerful simulation tools that make possible both the generation of raw scientific data and the manipulation of this information to enhance understanding.

The institutional nature of the initiative makes it possible for anyone with a good idea to perform teraops-scale (more than 1 trillion mathematical operations per second) parallel computing to achieve breakthrough scientific discoveries. First developed in the late 1980s, parallel computing attacks huge mathematical problems with a number of identical (and typically inexpensive and common) processors simultaneously sharing a computational task. Parallel computing differs from traditional supercomputing, in which a few high-cost, specially designed vector processors perform numerical calculations.

"Parallel computing brings formerly insoluble problems into solution range and makes formerly difficult-to-solve problems routine," McCoy says. Examples of parallel supercomputers include the Laboratory's ASCI supercomputers manufactured by IBM—the classified Blue Pacific (3.9 teraops) and White (12.3 teraops) and the unclassified Blue (0.7 teraops) and Frost (1.6 teraops)—and M&IC's

TeraCluster2000 supercomputer, built by Compaq, operating at 0.7 teraops.

McCoy explains that when the initiative was launched in the mid-1990s, many scientists, having experienced difficulties in securing sufficient and sustained access to unclassified centralized computing, had forsaken high-performance computing entirely. Instead, they had taken advantage of the opportunities presented by new and relatively inexpensive and powerful desktop computers in their offices or used terminals tied to scientific workstations owned by their research programs. He calls the situation at the time "a desktop diaspora."

While desktop workstations continue to be an important tool for scientific computing research, these machines do not always provide the necessary computational power. "Exclusive recourse to workstations, in the absence of access to the most powerful centralized systems available, could have left many Lawrence Livermore researchers unequipped with the computer tools they needed to remain competitive in the next decade," says McCoy.

"Until this time, most scientists were performing two-dimensional simulations because three-dimensional simulation was quite difficult, if not impossible, with the computers at their disposal," says Greg Tomaschke, M&IC computing systems leader. "Livermore has a long history of preeminence in computational simulation, and people were concerned they'd start to lose this advantage." M&IC has ensured that

On the Horizon: The Terascale Simulation Facility

Looking to the future, Lawrence Livermore managers are planning a new facility to house their newest supercomputers. Called the Terascale Simulation Facility, the building will cover 25,000 square meters and have offices for 288 people. A two-story main computer structure will provide the space, power, and cooling capabilities to simultaneously support two future supercomputer systems for the National Nuclear Security Administration's Accelerated Strategic Computing Initiative (ASCI). Nearby, the Livermore Computing Center will house the unclassified computing resources of the Multiprogrammatic and Institutional Computing Initiative.

The computer structure will be flanked on the south side by a four-story office building. The building will have space for Lawrence Livermore staff, vendors, and collaborators from the ASCI University Alliance and visiting scientists from Los Alamos and Sandia national laboratories.

Lawrence Livermore computer managers say that no existing computer facility at Livermore is adequate, nor can any be modified sufficiently, to accommodate the newest supercomputers, which are the largest computational platforms ever built. These platforms are so powerful—exceeding 50 teraops, or 50 trillion operations per second—that they are increasingly labeled ultracomputers. Current plans are to site a 60-teraops system in the Terascale Simulation Facility by 2004.

The Terascale Simulation Facility will house the largest supercomputers ever built and offices for nearly 300 people.

Livermore researchers are at the forefront of simulation science, he says.

Serving Programs and Individuals

The M&IC is so named because it serves both research programs (multiprogrammatic) and individual (institutional) researchers. A research program can either purchase a block of time on existing machines or, more popularly, share in the investment in new equipment, which gives the program access proportional to its investment.

The research program designates a principal investigator to select users of the resources it has paid for. The Livermore Computing Center furnishes the principal investigators with Internet-based records of the computing time the program is entitled to and how much of the time has been used, in total or by individual researcher. A hotline is available to answer technical computing questions and provide account information.

M&IC also grants computer time to individual researchers, independent of their connection to a program, whose work is often viewed as less mainstream than most efforts in their particular research specialties. The researchers are selected through a short proposal process by the Institutional Computing Executive Group (ICEG), composed of users from the Laboratory's research directorates. The ICEG provides general oversight of M&IC, in cooperation with the Computation associate director and the Laboratory's deputy director for Science and Technology.

"The ICEG is the most important Livermore Computer Center link into the user community," says McCoy. "The M&IC Initiative is based on the strong bonds of support and advice that exist between the ICEG and the center. As a result, Livermore supercomputing has become an institutional resource much like the library, a place where researchers from any program can expect resources to support their research."

Researchers can monitor the status of their simulation from their desktop computer. Those who sign up for long simulation runs can see their place in the queue. Once the simulation begins, they can open a window on their computer screen to watch the calculations. "We've made investments in customer (user) support that significantly exceed the norm for computing centers worldwide because we know the extraordinary challenges scientists face with writing and running large parallel calculations," says McCoy.

Providing Capacity and Capability

Under the M&IC Initiative, the Livermore Computing Center has acquired increasingly more powerful clusters or groups of computers. The computers are of two conceptually different kinds: capacity computers and capability computers. Capacity computing is designed to handle jobs that don't require a lot of computing horsepower or memory. It is often used for quick turnaround of a large number of small to moderately sized simulations.

Capability computing uses a substantial fraction of the entire computing power of a supercomputer to address a large-scale scientific simulation in three dimensions. "The driver for capability computing usually is the need for large amounts of memory, which means harnessing many processors to work together," says Tomaschke. A capability computing resource can only serve a few users simultaneously.

McCoy says that any effective computational environment is supported by a capacity foundation. "Capacity allows users to develop the applications and work the studies that are necessary to conceive of, develop, and debug capability applications," he says. M&IC managers, working with the ICEG, developed a strategy early on to first build a capacity foundation and keep it current. They then devised a second strategy to build access to capability computing, either through a partnership with the ASCI program or through unique research and development relationships with major vendors. Both strategies have proved effective.

The Compaq 8400 Compass Cluster was the first capacity computer resource sponsored by the M&IC Initiative. Delivered in 1996, the Compass Cluster consists of 8 nodes (rack-mounted computers), with each node possessing between 8 and 12 central processing units (CPUs, or microprocessors). In all, 80 CPUs provide 7 gigaops (7 billion calculations per second) of computing power, 56 gigabytes of memory, and about 900 gigabytes of disk space. A replacement for Compass, scheduled to arrive late this year, will provide a total of 192 gigaops.

To increase capacity computing, M&IC acquired TeraCluster in late 1998. In all, TeraCluster consists of 160 CPUs, 80 gigabytes of memory, and over 1.5 terabytes (trillion bytes) of disk space and provides about 182 gigaops of computing power. It is closely integrated with the Compass Cluster.

In September 2000, the Livermore Computing Center took delivery of the Linux Cluster, which generates 42 gigaops of computing power. It is composed of 16 advanced Compaq nodes, with each node having 2 CPUs and 2 gigabytes of memory. The machine is used to increase computing capacity and provide users with an opportunity to evaluate the potential advantages of the Linux operating system. McCoy says that Linux represents a radical departure because it is not a proprietary operating system. Success with the machine could lead to the procurement, next year, of Linux parallel processing systems using a high-performance interconnect and advanced Intel processors.

The M&IC Initiative has also acquired a system manufactured by Sun Microsystems called Sunbert, which provides 12 gigaops of power with 24 CPUs, 16 gigabytes of memory, and approximately 600 gigabytes of disk space. The system is designed to allow access by Livermore researchers who are foreign nationals from sensitive nations.

Capability Power

The M&IC's most important capability platform, the 680-gigaops TeraCluster2000 (TC2K) parallel supercomputer, arrived last year. The machine was the result of a three-year Cooperative Research and Development Agreement between computer scientists at Livermore and Compaq Computer Corporation to evaluate a new supercomputer design based on Compaq's 64-bit Alpha microprocessor and Quadrics Corporation's interconnects and software. The alliance resulted in the Compaq SC Series of supercomputers, of which TC2K is serial number 1.

Tomaschke says that the most important aspect of Livermore's role in the collaboration with Compaq was providing advice based on many years of experience with supercomputers. "We gave Compaq important feedback about what scientists require for doing three-dimensional simulations, such as operating enormous file systems."

The TC2K consists of 128 nodes, with 4 Alpha processors per node. In total, the machine has 512 CPUs, 256 gigabytes of memory, and 10 terabytes of disk space. The 128 nodes are partitioned like a giant hard disk. The largest partition is dedicated to the most complex simulations, while a small partition permits a researcher to interact with the machine in real time. Occasionally, all nodes are freed up for a single task, such as experiments to determine if a code will scale properly when the number of nodes increases sharply.

Limited availability of TC2K began early this year, with 25 projects (involving about 100 researchers) "shaking down" the machine. It became generally available in August, vastly increasing computing capability to all unclassified researchers.

TC2K represents one of three capability resources. The second resource is ASCI Blue, the 740-gigaops unclassified portion of the ASCI Blue Pacific system, which has 282 nodes with 4 IBM Power PC processors each. The third resource is ASCI Frost, the unclassified version of ASCI White. This system features 68 nodes with 16 powerful 1.5-gigaops IBM processors each and 16 gigabytes of total memory. This computer peaks at 1.6 teraops and is both the most modern and most powerful unclassified computer on site. Although both Blue and Frost are primarily dedicated to the ASCI mission, significant access has been made available to a number of Livermore science teams.

TC2K's capability, combined with that of ASCI Blue and ASCI Frost, provides unprecedented unclassified computing capability for a national laboratory, says McCoy. Researchers perform code development and limited simulations on capacity machines, complex three-dimensional simulations

The TeraCluster2000 (TC2K) Compaq supercomputer, operating at 0.7 teraops (trillion operations per second), has 512 central processing units arranged in 128 nodes, 256 gigabytes of memory, and 10 terabytes of disk space.

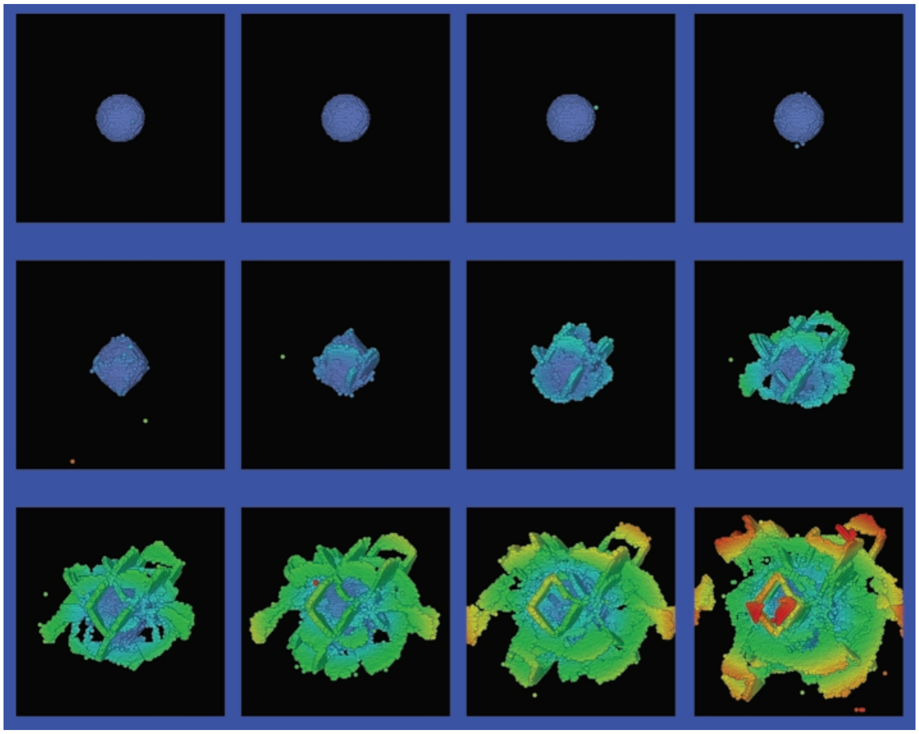

A sequence of snapshots from a simulation showing how a growing fracture deforms a stressed copper crystal at the level of atoms. The first row shows no defects. The second row shows defects forming on the surface. The third row shows defects shooting off from the surface in a process that forms a large hole.

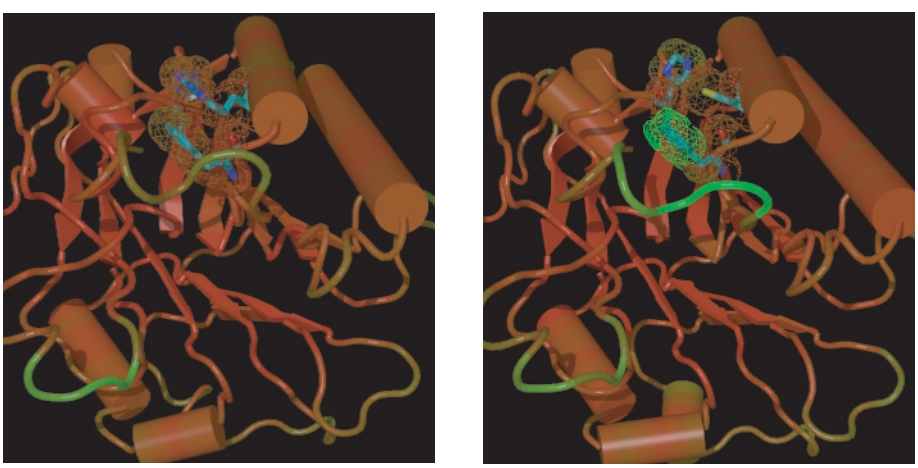

The TC2K supercomputer is used to gain insight into the human Ape1 enzyme, a protein that repairs DNA. The simulation at left is a healthy protein. The simulation at right is a version of the protein that contains a single amino acid substitution. This variant shows much more motion in the front loop of the protein; the motion is a means of recognizing DNA damage. The coloring indicates the amount of intramolecular motion, with reddish-brown being the least motion and greenish-blue the most motion.

on TC2K and Blue, and the most demanding runs on Frost. Access to the limited unclassified ASCI resources is extremely competitive.

Never Before Attempted

The resources provided by the M&IC Initiative are permitting researchers to generate simulations that, in many cases, were never before attempted for lack of computing power. As a result, the Laboratory is at the forefront of simulating a wide range of physical phenomena, including the fundamental properties of materials, complex environmental processes, biological systems, and the evolution of stars and galaxies.

For example, physicist Burkhard Militzer of the Quantum Simulation Group in the Physics and Advanced Technologies Directorate is using TC2K to simulate how gases such as hydrogen and oxygen behave under extreme pressure and to compare those simulations with results of gas-gun experiments done at Livermore. Militzer uses JEEP, a parallel supercomputer code developed by Livermore physicist Francois Gygi. The simulations typically require upward of 300 hours of processing time. Because there is a time limit of 12 continuous hours on TC2K, the simulations run in chunks.

Militzer says that he would like to use JEEP's quantum mechanics capability to simulate the weak hydrogen bonds that keep two DNA strands in their helix. "It's very difficult to do accurately because of all the water molecules surrounding the DNA," he says. "TC2K enables a new class of projects. I wouldn't even begin to think about running a simulation of DNA hydrogen bonding without having TC2K."

Examining the Birth of Cracks

Physicists Robert Rudd and Jim Belak run simulations on the Compass

Cluster. The simulations examine in microscopic detail the birth of fractures in metals such as copper under the extreme stresses of a shock wave. The molecular dynamics simulations are done in a nanometer-scale box holding about one million virtual atoms.

The simulations, actually a sequence of snapshots from a movie, depict a copper crystal that is deformed by the growing fracture over a period of 60 picoseconds (trillionths of a second). Only the atoms at the fracture surface or in crystal defects are shown. The defects, known as dislocations, can be seen shooting off in a process that forms an increasingly large hole or fracture. Rudd says, "We're interested in learning more about how voids grow and how the material deforms around them."

"We have great confidence in our simulations," adds Rudd. He is planning to use ASCI Blue to vastly expand the length of the simulated piece of metal and to simulate much longer time periods.

Lawrence Livermore biological scientists in the Computational Biology Group have been one of the most visible users of the M&IC Initiative. (See *S&TR*, April 2001, pp. 4–11.) The researchers have produced stunning depictions of DNA and proteins that reveal the exact mechanisms of key biological processes. Parallel supercomputers are ideal for this kind of simulation because they excel at modeling the interactions of large numbers of atoms contained within biological macromolecules.

Researcher Daniel Barsky of the Computational Biology Group has been studying the dynamics of Ape1, an enzyme responsible for repairing a common form of DNA damage called abasic lesions. Barsky uses the TC2K to compare the degree of intramolecular motion of normal Ape1 with a variant found in which the enzyme contains a single amino acid substitution.

Computer scientist and geophysicist Shawn Larsen has been using TC2K to do three-dimensional simulations of oil exploration problems and seismic wave propagation. Larsen's E3D code is optimized for parallel supercomputers. (See *S&TR*, November 2000, pp. 7–8.) His oil exploration modeling entails so-called elastic simulations that are about 250 times more computationally intensive than standard acoustic (sound waves in air) simulations. Elastic simulations provide better details of subsurface geology, which is essential to oil exploration efforts, by depicting both the S (shear) waves and the P (compressional) waves that travel in water and earth. With computers such as TC2K, realistic three-dimensional elastic simulations are now possible.

Climate Simulations Set the Pace

Livermore researchers in atmospheric sciences were among the first to take advantage of M&IC's new capabilities. Having performed parallel computing for 10 years, they had already achieved some of the most advanced simulations ever done. They had also made their simulation codes portable, that is, easily adaptable to different computers. As a result, it did not take long to adapt their codes to TC2K and ASCI Frost.

They put TC2K to the test to study the effectiveness of ocean carbon sequestration, a proposed approach for mitigating global warming. In one sequestration model, the carbon dioxide generated by industrial operations would be injected into the oceans instead of being emitted into the

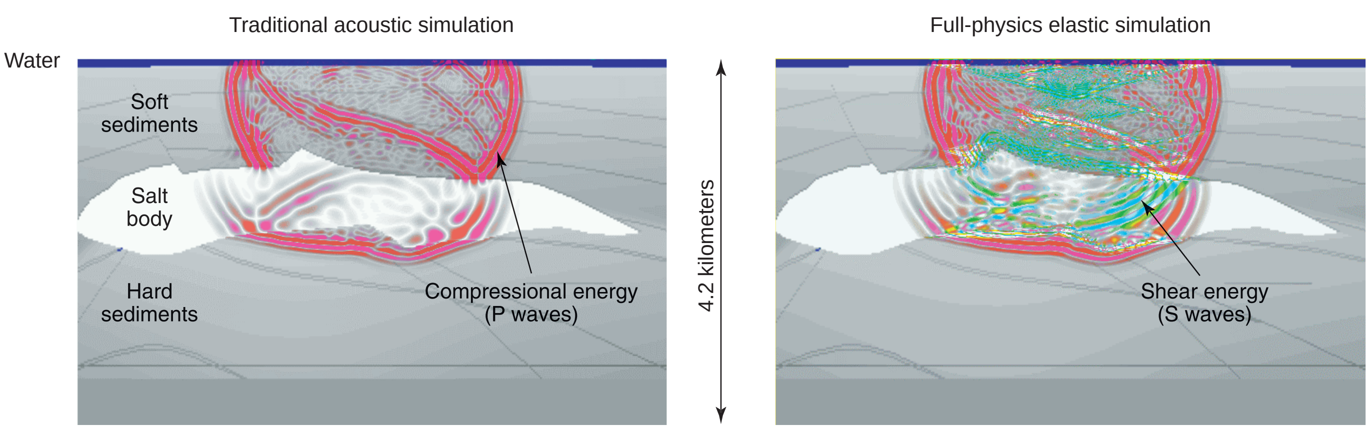

The TC2K supercomputer is able to perform three-dimensional simulations of seismic wave propagation with Livermore's E3D code that is optimized for parallel supercomputers. The image on the left used 1 central processing unit (CPU) and 0.3 gigabyte of computer memory in a traditional acoustic simulation of an underwater deposit. The image on the right used 240 CPUs running for up to 18 hours and 85 gigabytes of computer memory to generate a full-physics elastic simulation. The image contains a great deal more detail, including seismic S (shear) waves that travel in the earth.

atmosphere. However, some of the injected carbon dioxide would eventually leak into the atmosphere, where it would contribute to climate change.

To evaluate this approach to mitigating global warming, the Department of Energy formed a Center for Research on Ocean Carbon Sequestration, located at Lawrence Livermore and Lawrence Berkeley national laboratories. For one of the center's studies, Ken Caldeira, codirector of the center, along with colleagues Philip Duffy of Atmospheric Sciences and Mike Wickett of the Center for Applied Scientific Computing, used TC2K to evaluate the effectiveness of ocean carbon sequestration over a period of 100 years. "We want to know how much of the injected carbon dioxide would leak out of the ocean and at what rate," says Duffy. Their simulations showed that leakage into the atmosphere is much less when carbon dioxide injection is done at greater depths. The simulations also showed that maximum leakage may occur far from the injection site in areas where there is vigorous overturning of the ocean, such as in the North Atlantic. The simulations, the highest resolution of their kind ever done, used up to a hundred TC2K nodes and required a total of about 10,000 CPU hours.

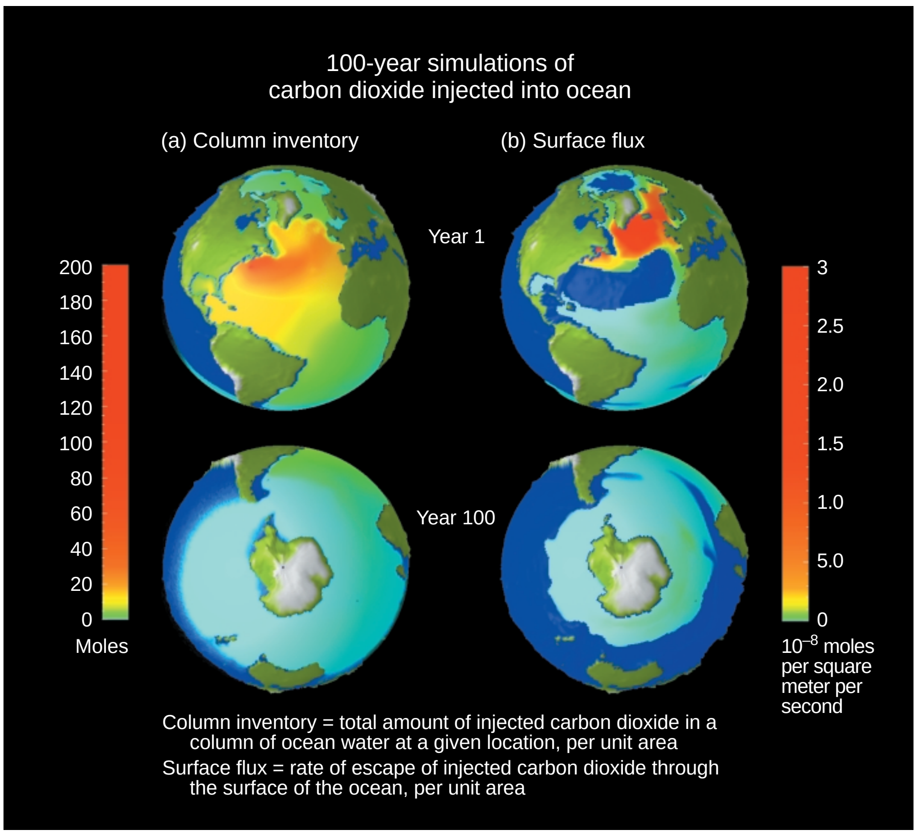

Multiprogrammatic and Institutional Computing supercomputers are making it possible to evaluate the potential effectiveness of injecting carbon dioxide into the ocean to mitigate global warming. These 100-year simulations depict what happens to carbon dioxide injected into the ocean off New York City at a depth of 710 meters. The column inventory, shown in the two spheres at the left, depicts the carbon dioxide traveling a great distance in 100 years. The surface flux, at right, shows that maximum escape of the injected carbon dioxide can occur far from the injection site in areas where there is vigorous overturning of the ocean, such as the North Atlantic.

Duffy and other atmospheric researchers, together with collaborators from Livermore's Center for Applied Scientific Computing, have used TC2K and other DOE supercomputers to perform the highest-resolution global climate simulations ever done. Duffy notes that as the simulation resolution gets finer, topographic features like mountains and valleys are represented better; these features have a direct bearing on weather. The simulations show that in some regions, maximum warming will occur in high-elevation regions because of a snow-albedo feedback: warming causes reduced snow cover, which in turn amplifies the warming by reflecting less sunlight back into space.

The simulations were performed in part on TC2K, in part at the National Energy Research Supercomputing Center at Lawrence Berkeley, and in part on ASCI Frost. "TC2K has really allowed us to push the limits of model resolution," says Duffy. "We're doing things that no other researchers can duplicate."

Unprecedented Ozone Studies

Atmospheric scientists in the Atmospheric Chemistry Group have used a hundred processors for a total of about 400 hours in their studies of the atmosphere. They have the only atmospheric chemistry model, IMPACT, that is capable of simulating the chemical reactions occurring in both the troposphere (the 10 kilometers of the atmosphere closest to Earth) and stratosphere. Past studies modeled the troposphere and stratosphere independently because of computational limitations. TC2K has enough computational power to permit coupling the troposphere and stratosphere in an atmospheric chemistry simulation capable of accurately predicting ozone concentrations. The results showed that studying interactions between these two regions of the atmosphere is important for understanding global and regional ozone distributions.

Understanding the ozone distribution throughout the atmosphere is crucial to the ability to predict not only the possibility of future stratospheric ozone depletion but also Earth's radiation balance and the magnitude of global warming. The scientists' model included resolutions of 2 degrees latitude by 2.5 degrees longitude, 46 levels of altitude from Earth's surface to 60 kilometers, tropospheric and stratospheric chemistry and physics involving 100 chemical species and 300 chemical reactions, and weather dynamics. "Because our codes are CPU-intensive, they do well on TC2K," says atmospheric scientist Doug Rotman. "The machine excels at big problems involving a lot of parameters." He says that while one goal is to keep increasing the resolution of the simulations, another goal is to include additional physics to make simulations more realistic.

Engineers David Clague, Elizabeth Wheeler, and Todd Weisgraber and University of California at Berkeley student Gary Hon have been using TC2K to perform three-dimensional simulations of both the stationary and mobile particles in portable microfluidic devices. These devices are being designed by Livermore researchers to automatically detect and identify viruses, bacteria, and toxic

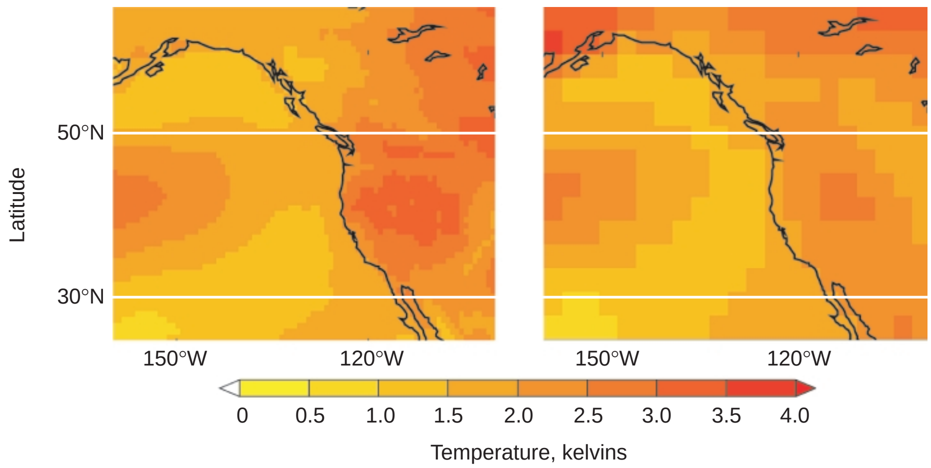

Simulations depict changes of surface air temperature between 2000 and 2100 in the western U.S. The simulation on the left, produced on Livermore's ASCI Frost and TC2K supercomputers, has a 75-kilometer resolution. The simulation on the right, produced on the National Energy Research Supercomputing Center at Lawrence Berkeley National Laboratory, has a 300-kilometer resolution. Its finer resolution gives a more detailed prediction that should be more accurate.

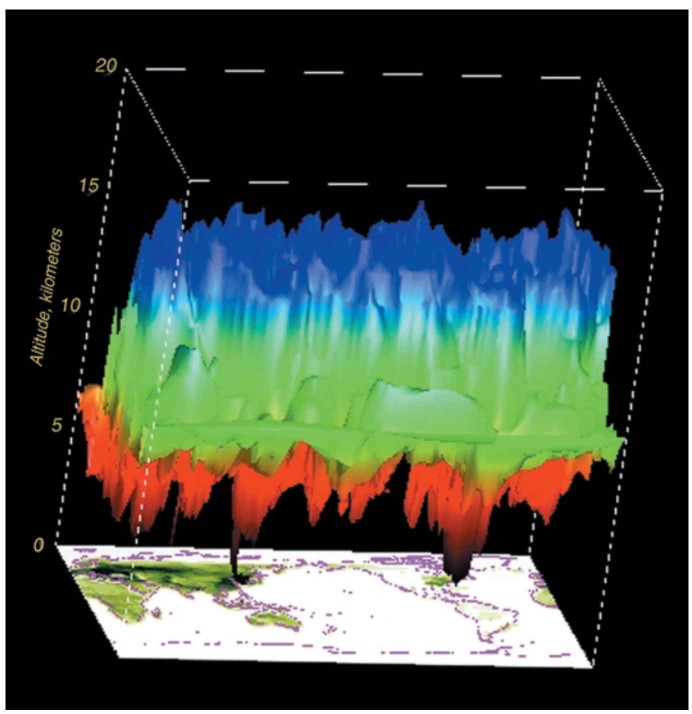

The TC2K supercomputer permits simulations that for the first time couple the troposphere and stratosphere in modeling the distribution of ozone. Modeling results show that in various locations, ozone in the lower stratosphere, which is typically in 100-parts-per-billion concentrations, is transported to near-surface altitudes.

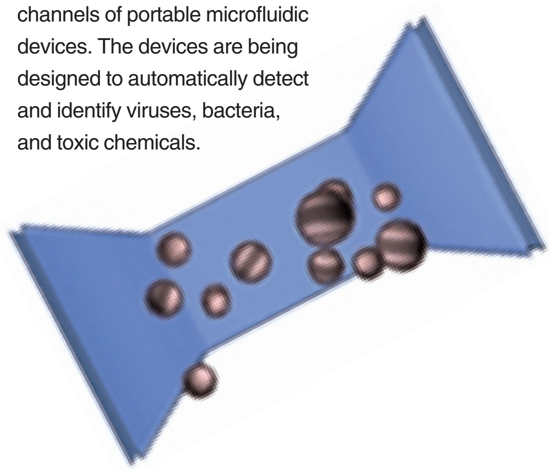

This simulation using TC2K shows the influence of electric fields on molecules traveling in narrow channels of portable microfluidic devices. The devices are being designed to automatically detect and identify viruses, bacteria, and toxic chemicals.

chemicals. (See *S&TR*, November 1999, pp. 10–16.) The devices have channels from 20 to 200 micrometers deep and up to a millimeter wide through which fluids travel. Because the channels are so small, intermolecular forces, which are typically masked in laboratory-scale instruments, affect the behavior of particles. The simulations show how beads and macromolecules are affected by each other's electric fields as they travel through a channel.

Initiative a Big Success

McCoy is pleased that M&IC computing resources have been so well received. One sign of the initiative's success has been the growing competition for the finite resources and the occasional wait of several days to begin big simulation runs.

"The desktop diaspora is over," says McCoy, and the result is unprecedented simulations and outstanding science. "We have achieved a balance in understanding what can be best done on the desktop and what can be best done in the experimental computational facility. We have developed close partnerships with our science teams, and we are already planning the next steps."

McCoy adds it is the "momentum based on continual change that keeps me and the Scientific Computing and Communications Department engaged and interested." That interest, he says, is based in part on the newest generation of simulations. They promise significant discoveries in science as Livermore researchers continue to elevate simulation to a level equal to that of theory and experiment.

—*Arnie Heller*

Key Words: Accelerated Strategic Computing Initiative (ASCI), ASCI Blue Pacific, ASCI Frost, ASCI White, atmospheric sciences, carbon sequestration, Compass Cluster, elastic simulations, E3D, Institutional Computing Executive Group, JEEP, microfluidic devices, Multiprogrammatic and Institutional Computing Initiative, parallel computing, Scientific Computing and Communications Department, Sunbert, supercomputers, TeraCluster, TeraCluster2000 (TC2K), Terascale Simulation Facility.

For further information contact Mike McCoy (925) 422-4021 (mccoy2@llnl.gov) or Greg Tomaschke (925) 423-0561 (tomaschke1@llnl.gov).

About the Scientist

MIKE McCOY, deputy associate director of Scientific Computing and Communications in Livermore's Computation Directorate, received his A.B. (1969) and Ph.D. (1975) in mathematics from the University of California at Berkeley. He joined the Laboratory in 1975 as a student employee. Upon completing his doctoral dissertation, he became a staff scientist in the National Energy Research Supercomputer Center, where he took responsibility for the development of algorithms for plasma codes. He went on to become group leader of the center's Massively Parallel Computing Group and then its deputy director. In the latter role, he directed the procurement of the 256-processor T3D computer, which was used for unclassified science. Now, as deputy associate director for Computation, he continues to support programs to advance network communications and security and to foster the development and integration of systems of powerful computers. He has worked with Livermore science teams to establish institutional computing to provide Laboratory scientists with access to powerful simulation environments. This sharing of institutional resources is part of his vision for enhancing the role of simulation in the scientific triad of theory, experiment, and simulation.

Further Developments in Ultrashort-Pulse Lasers

LIVERMORE continues to push the frontiers of laser science. Researchers have taken the ultrashort pulses of the record-shattering Petawatt laser and found new uses for them. They are also applying a novel technology to create extremely short laser pulses with high average power and high energy to access and investigate extreme-field conditions. The technology improvements will benefit the Stockpile Stewardship Program as well as national defense and manufacturing.

The Petawatt laser operated for three years and routinely produced more than 500-joule laser pulses lasting 500 femtoseconds—less than a trillionth of a second. Experiments with the Petawatt evaluated the fast ignitor method of achieving inertial confinement fusion, generated powerful electrons or x rays for radiography research, and produced short, powerful gamma rays for nuclear physics experiments. (See *S&TR*, March 2000, pp. 4–12.) In addition, the discovery of intense, high-energy, collimated proton beams emitted from the rear surface of Petawatt laser targets has opened the way to new applications such as proton radiography. The Petawatt laser still holds the world's record for the highest peak power ever achieved by a laser.

The Petawatt operated on one of the 10 beam lines of Livermore's Nova laser. When the Nova laser was decommissioned in 1999, the Petawatt went with it. But work on short-pulse lasers by no means stopped, notes physicist Mark Hermann, associate program leader for Livermore's Short-Pulse Lasers, Applications, and Technology program, known as SPLAT, which is a part of the National Ignition Facility (NIF) Programs Directorate. His team of about 30 people is advancing the science of short-pulse lasers and applications, developing new laser components, fielding advanced laser systems, and developing new optical components and optical fabrication technologies. There is also

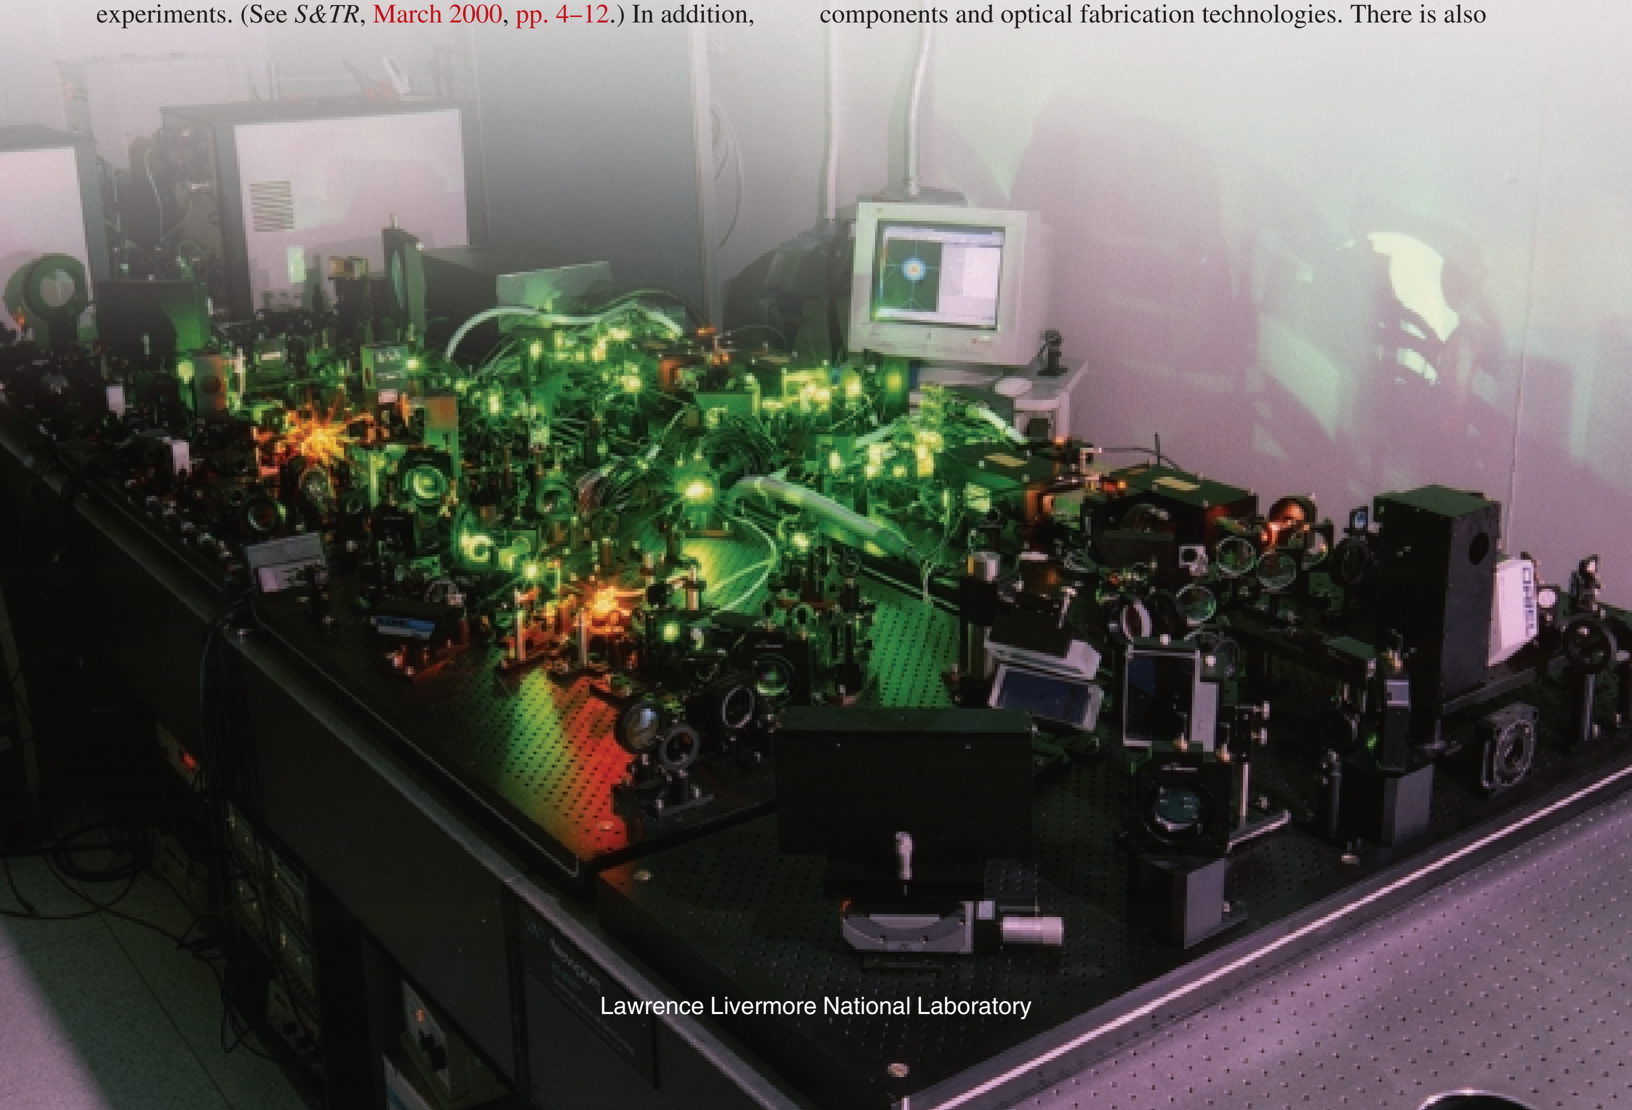

an active program in short-pulse technology in the Physics and Advanced Technology (PAT) Directorate. This research stems from the need to develop high-temperature plasma sources and accurate plasma probes for high-energy-density materials research.

In SPLAT, a diverse set of challenging projects focused on developing high-average-power, short-pulse lasers for a variety of customers is under way. Current SPLAT-developed laser systems use conventional titanium-doped sapphire (Ti:sapphire) amplifiers, but now the team is developing new chirped-pulse amplifier technologies geared toward high average power. One is a direct, diode-pumped, chirped-pulse amplifier laser crystal that promises efficient, compact, and robust picosecond-pulse laser systems. Another is an optical-parametric chirped-pulse amplification (OPCPA) technique, described in more detail below.

The SPLAT team is using a short-pulse laser to create unique nanocrystals and gain knowledge about the novel properties of nanostructures. This knowledge affects the basic sciences, from solid-state physics to biology. Being able to synthesize nanocrystals of specific size and properties, at an industrial rate, may revolutionize the field of nanotechnology and enable a broad sector of manufacturing, from semiconductors to pharmacology.

The team is collaborating on a project that integrates a short-pulse laser with a Livermore linear accelerator for stockpile stewardship applications. Supporting the team's efforts, Livermore's Diffractive Optics Group is developing new optical technologies and fabricating new optics for petawatt-class lasers around the world, for the National Ignition Facility, and for the National Aeronautics and Space Administration to use in space-based telescopes. The group currently produces the world's largest diffraction gratings.

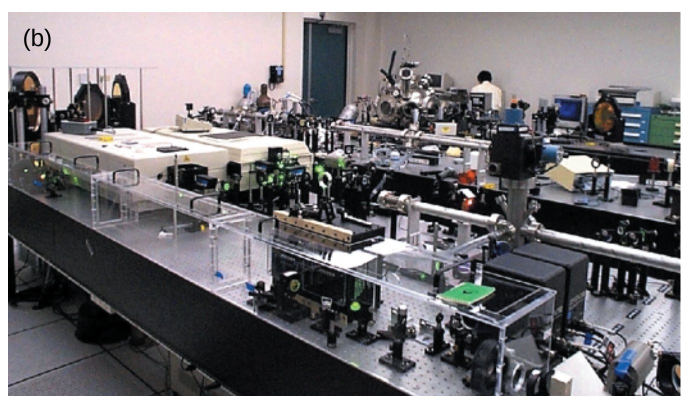

(a) The100-megaelectronvolt linear accelerator and (b) the Falcon ultrashort-pulse laser are being integrated to produce a short-pulse x-ray source. Livermore scientists will use the x rays to probe the dynamics of materials under shock conditions.

New Path to a Short Pulse

An interesting new development in short-pulse laser technology has been the emergence of OPCPA. Laser-pumped nonlinear crystals made of beta-barium borate (BBO) would replace the Ti:sapphire used in the Petawatt and other conventional lasers as the preamplifier. In a Ti:sapphire regenerative power amplifier, a pulse passes 10 to 100 times through a regenerative cavity, increasing in energy with each pass. By the time it leaves the amplifier, its energy has increased by 10 million, from about 1 nanojoule to approximately 10 millijoules. In contrast, a pair of BBO crystals can produce the same energy gain with a single pass of the light pulse.

With a regenerative amplifier, a tiny bit of energy leaks out with each round trip of the laser pulse. If this leak, or prepulse, is not attenuated, it may cause a preplasma, which changes the coupling of the laser to the target. Many stockpile stewardship experiments use lasers to probe materials essential to nuclear weapons to learn more about their behavior. Keep in mind that the prepulse causes changes that are miniscule by most standards, but when powerful laser pulses of less than a trillionth of a second are used to study detailed physics processes, even minor changes can be significant.

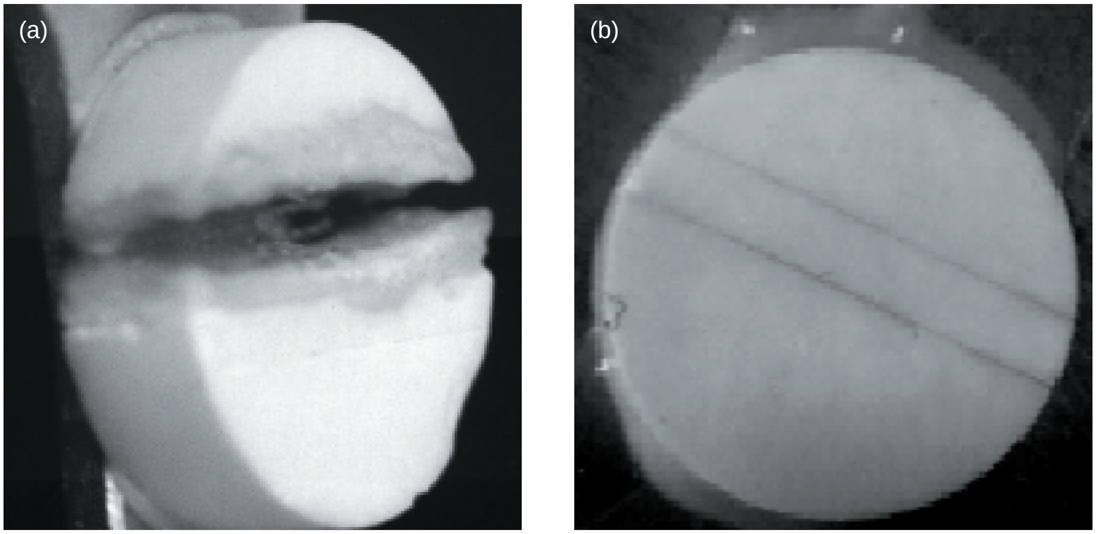

(a) A laser with 500-picosecond pulses caused the high explosive LX-16 to burn during cutting. (b) In Livermore's system, the 150-femtosecond laser pulses are so short and fast that they deliver virtually no heat to the area being cut.

Livermore did not invent OPCPA. But, according to Hermann, "Livermore is pushing the frontier of OPCPA technology, combining Livermore's unique expertise in high-beam-quality, high-average-power lasers and nonlinear optics."

A major application for OPCPA will likely be in laser machining, which requires high average power, high beam quality, and ultrashort pulses (about 20 to 1,000 femtoseconds). Unlike other chirped-pulse amplification approaches, OPCPA produces negligible thermal aberrations that in turn cause degradation of the laser beam. Although not yet demonstrated at high average powers, an OPCPA laser should be able to produce hundreds of watts or even kilowatts of average power with high beam quality. In contrast, most conventional Ti:sapphire chirped-pulse lasers, including Livermore's systems, have operated at 20 watts or less. Higher power should translate into faster production and better process control during machining.

Much of the SPLAT program's work is related to stockpile stewardship and improving Livermore's ability to verify the safety and reliability of the nation's aging nuclear weapons stockpile. Before the arrival of OPCPA technology, the team built the Falcon, a 3-terawatt, 35-femtosecond laser facility with a moderate repetition rate, to use as a material probe. A joint team of SPLAT and PAT personnel recently began to integrate the output from Falcon with the electron beam generated by Livermore's 100-megaelectronvolt linear accelerator. Together, the laser and the accelerator will be an advanced light source whose ultrafast and ultrabright pulsed x rays will be used as probes for dynamic studies of solid-state and chemical systems.

Hermann notes that the Falcon and other short-pulse lasers at Livermore may benefit from being upgraded with the OPCPA system. "Controlling and in some cases eliminating the prepulse is desirable," says Hermann. "It means that researchers will be able to control the experimental initial conditions of the laser material dynamics."

Benefits Abound

Emerging technologies for optical-parametric chirped-pulse amplification and diffractive optics will soon find their way into several Livermore lasers, from small high-average-power systems for manufacturing to high-energy systems such as NIF, the 192-beam laser being built to support stockpile stewardship science research. NIF and SPLAT personnel are assessing the potential for adding these technologies to produce short pulses on NIF. Combining ultrashort pulses with the powerful, multimegajoule capacity of NIF would result in a unique system, one that may be able to demonstrate fast ignition for the production of fusion energy, increase NIF's stockpile stewardship capabilities, and investigate new areas of extreme-field science.

—Katie Walter

Key Words: diffraction gratings, Falcon laser, femtosecond laser machining and cutting, high-average-power lasers, nanocrystals, optical-parametric chirped-pulse amplification (OPCPA), Petawatt, ultrashort-pulse lasers.

For further information contact Mark Hermann (925) 423-8672 (hermann1@llnl.gov).

Simulating How the Wind Blows

ANALYZING and simulating atmospheric transport and dispersion is the stock in trade of Livermore's National Atmospheric Release Advisory Center (NARAC). In 1986, NARAC took on its first big test and successfully predicted the path that radioactive contaminants from the Chernobyl nuclear plant explosion would take across the USSR and Europe. Since then, NARAC has analyzed dozens of other chemical and nuclear accidents around the world. Understanding wind flow and turbulence in the atmosphere is at the heart of being able to predict the transport and dispersion patterns of atmospheric releases.

Those complex analyses were on a large regional or even global scale. Understanding how particles or gases move on a smaller scale is more difficult. The movement of a gentle breeze, for example, is not as simple as it looks. Trees, buildings, walls, fences, and the like cause eddies and other changes in the direction and speed of the breeze. A single building can create astonishing changes in flow patterns.

These small-scale effects matter for atmospheric scientist Bob Lee and colleagues Ron Calhoun, Steve Chan, John Leone, and Dave Stevens. They form a team that is seeking to understand and predict the behavior of biological and chemical agents released into the atmosphere within urban areas. Their work is part of Livermore's contribution to the Department of Energy's Chemical and Biological National Security Program.

Modeling and analysis of airflow and dispersion in urban settings may be done on three different scales, each of which can "see," or resolve, details at varying levels. The domains of the three scales are nested, with the highest spatial resolution at the smallest scale. At that smallest scale, the highly resolved simulation is of a single building or a few buildings, with a domain size of a few square kilometers and a time scale of a few minutes. The mid-range scale involves many buildings, tens of kilometers, and a dispersion time of one or two hours. These simulations can identify only clusters of buildings rather than individual buildings. The largest scale is urban/regional, which encompasses an entire urban and suburban area. Its size, as large as several hundred square kilometers, corresponds to the area of a typical regional weather forecast; the dispersion times may be many hours.

Lee and his colleagues have developed a finite-element-based computational fluid dynamics (CFD) modeling code known as FEM3MP. The code can handle either buoyant or heavier-than-air releases as well as processes involved in aerosol physics, bioagent viability degradation due to ultraviolet light, and surface heating and shading. FEM3MP treats buildings explicitly and incorporates parameters to

Developing atmospheric models for an urban setting requires taking many flow patterns into consideration. Flows around buildings are complex, with separation and stagnation zones, turbulent wakes, and vortices.

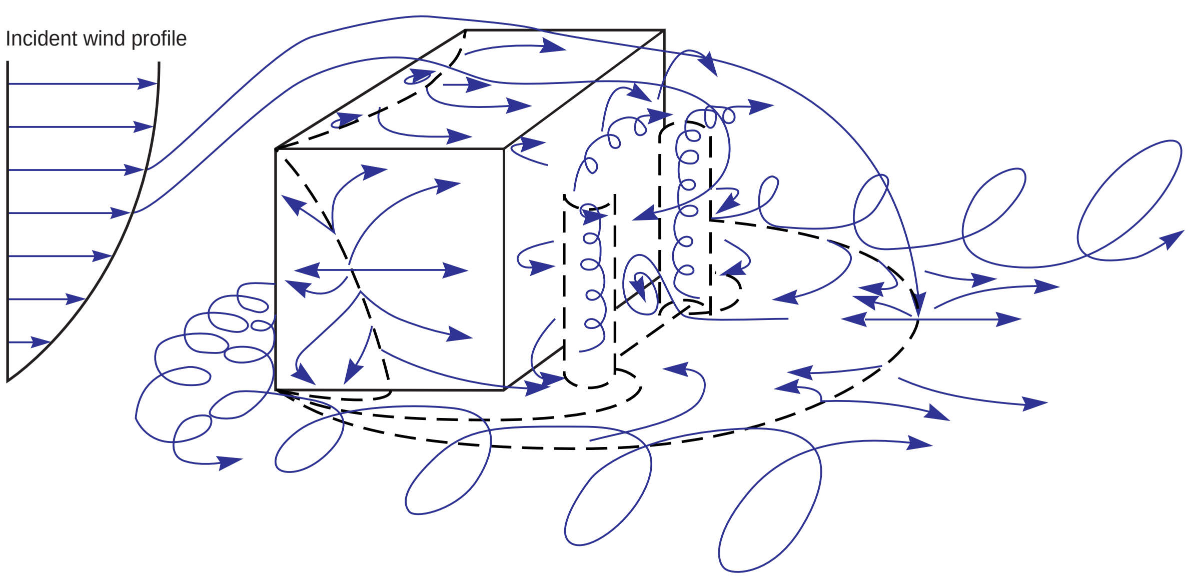

account for tree canopies. The simulations generated by FEM3MP represent turbulent air flow and the influence of atmospheric stratification on dispersion patterns within urban and surrounding areas.

Huge computers, especially the massively parallel machines of the Accelerated Strategic Computing Initiative (ASCI), have made it possible for models to incorporate a high level of detail and yet process calculations quickly. “We have been modeling atmospheric flow and dispersion since the early 1980s,” says Lee. “Much of our work has been in emergency response planning, which until recently had to rely on simple models. Not any more.”

Testing the model in the real world to assure its accuracy is as important as developing the model in the first place. Validation studies for urban air flow and dispersion models can be done in a wind tunnel as well as in the field, with real buildings. The first validation experiments for FEM3MP were conducted in a wind tunnel using blocks to represent buildings. The model simulations reproduced many important flow features such as separation zones, recirculation cavities between buildings, and interacting turbulent wakes. Only a sophisticated CFD model, such as FEM3MP, can capture these detailed flow features with sufficient accuracy.

Just One Building

Members of Lee’s team didn’t have to look far to test their model on a single building. They chose Livermore’s Building 170, home of NARAC, where most of them have offices. First, they ran FEM3MP to predict wind patterns around the building. This information told the field team where to place their instrumentation, consisting of about a dozen wind sensors that they moved to various locations around the building over a period of several months. They also made a few releases of an inert, harmless, odorless gas and tracked its evolution with remote sensing equipment and in situ instrumentation.

The team’s databank of experimental information from wind sensors and gas measurements then supplied the basis for a much improved model simulation—the “postexperiment” model simulation—of air flow around Building 170. Long-term meteorological data showed that the prevailing mean wind during the summer when the experiments were performed was out of the southwest. The top figure at the right compares the preexperiment and postexperiment simulations with actual field data for this case.

Lee and his colleagues found their models to be remarkably accurate, to within 10 percent for wind directions and

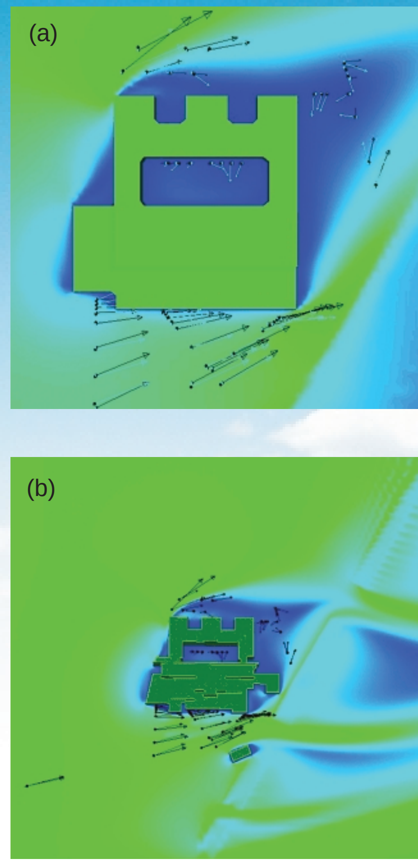

Researchers improved their air dispersion models when they incorporated field experimental data on wind patterns to calculate the consequences of gaseous releases into the models. The difference in simulation results is apparent in a comparison of (a) the preexperiment simulation with (b) the postexperiment simulation, which is much more complex and more accurately represents the details of the building.

Wind sensors were set up at various locations around Building 170 at Livermore to collect information about wind patterns.

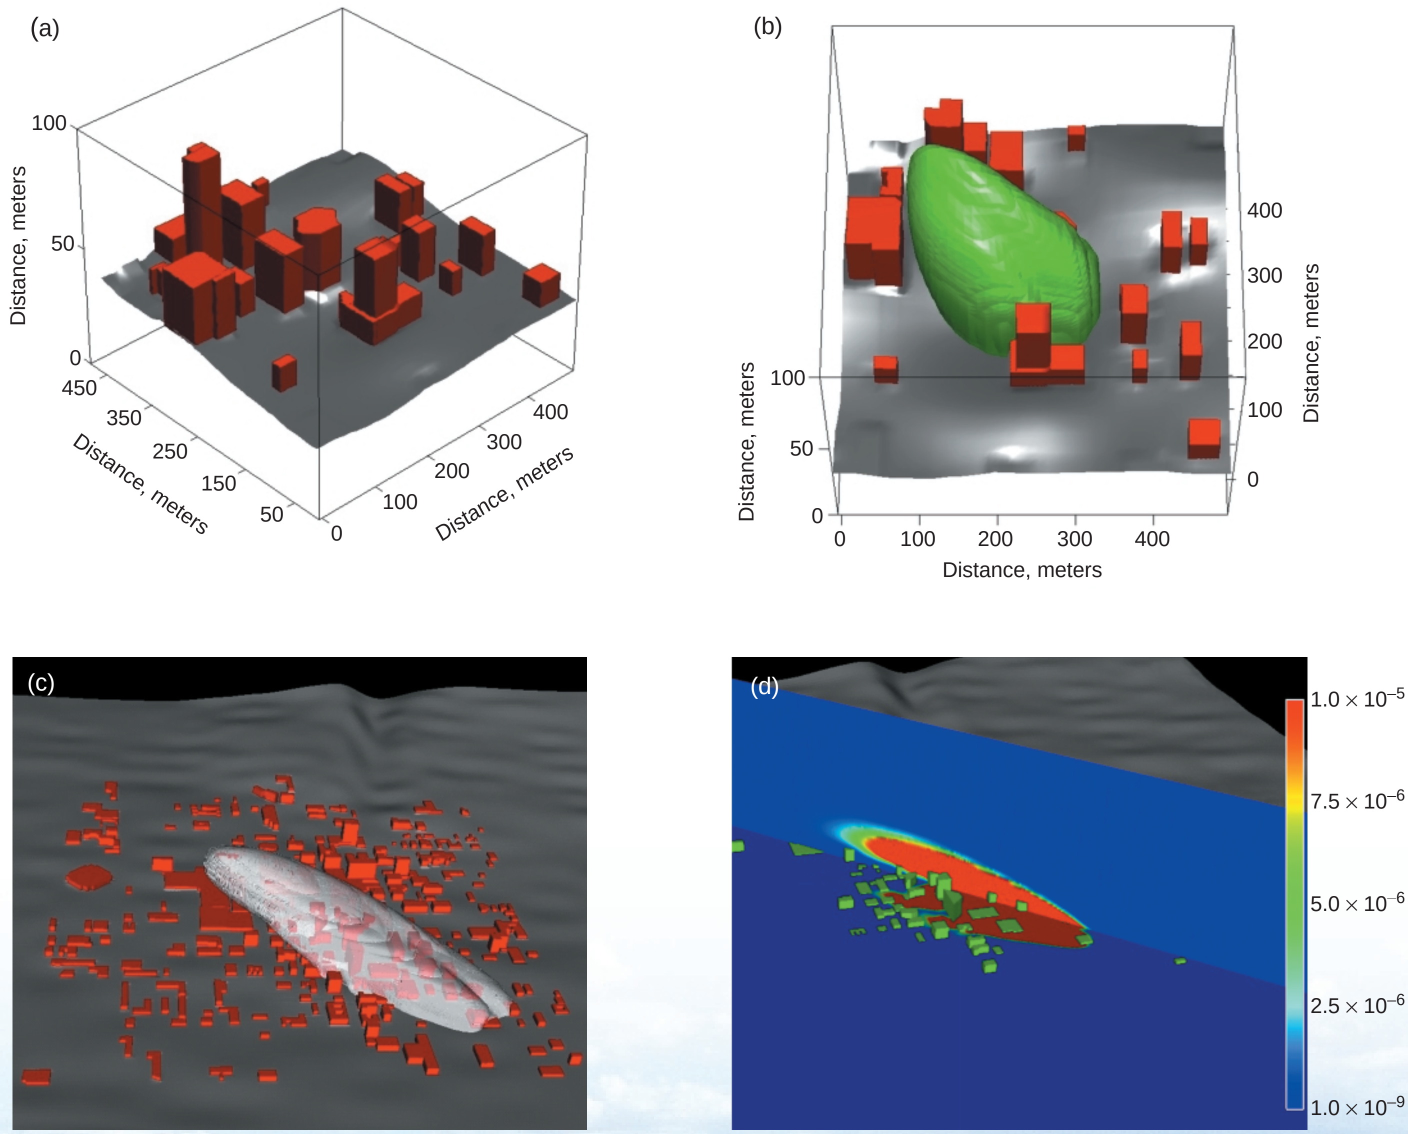

One simulation of the surface release of a generic tracer gas in downtown Salt Lake City. (a) Graph of the buildings within a 500-meter by 500-meter area near the source of the release. (b) An aerial view of the tracer plume over the greater downtown area 400 seconds after the release. (c, d) Cross-sectional views of the same plume showing the horizontal and vertical spread for the range of concentration depicted.

15 percent for wind speeds. The models successfully simulated many detailed flow features, such as the shedding of vortices from corners of the building and the blockage effect caused by a nearby row of 9-meter-tall eucalyptus trees. One of the most interesting simulations was one that used measured, second-by-second upwind data to "drive" the model calculation. This combination of experimental data and modeling represents a first step in the fusion of model and field data.

An Urban Scene and Beyond

Late last year, DOE completed an ambitious field experiment of atmospheric flow and dispersion in Salt Lake City, Utah. Livermore was a leader and major participant in the experiments and provided model simulations to the field team so that it would know where to place instrumentation.

Several DOE organizations, other federal agencies, and universities collaboratively collected field data relevant to all three scales of modeling, from a single building to the Salt Lake City basin to the entire region. Livermore used FEM3MP to generate flow and dispersion results for both the building scale and the basin scale. COAMPS (Coupled Ocean/Atmosphere Mesoscale Prediction System), a code developed by the Naval Research Laboratory and adapted by NARAC, provided forecasts for the regional scale. For the building scale, modeling focused on the effects of individual buildings or clusters of buildings. The basin scale modeling focused on terrain and thermal effects, in particular the air flows from the mountains east of Salt Lake City. The largest scale focused on regional-scale meteorology, within an area that includes parts of Idaho and Wyoming.

All of the simulations were massive, some with up to 100 million grid points. This work—simulating small-scale weather patterns and the dispersal of a tracer gas in the Salt Lake City area—especially benefited from the full computational power of the ASCI White computer. The building-scale exercises modeled up to 500 buildings to produce state-of-the-art, high-resolution simulations of air flow and dispersion. The figure on p. 18 shows views of some of the simulations.

"The Department of Energy is making a sizable investment in the computational-fluid-dynamics approach to urban dispersion modeling," says Lee. DOE's goal is to develop validated, multiscale computational models that support emergency preparedness, response, and detection of biological and chemical releases in urban areas.

—Katie Walter

Key Words: Accelerated Strategic Computing Initiative (ASCI), atmospheric dispersion, biological and chemical agents, Chemical and Biological National Security Program, computational fluid dynamics (CFD), National Atmospheric Release and Advisory Capability (NARAC), urban dispersion modeling.

For further information contact
Robert L. Lee (925) 422-1859 (lee34@llnl.gov).

Remembering E. O. Lawrence

"For those who had the good fortune to be close to him both personally and scientifically, he will always seem a giant among men."—*Luis Alvarez*

SCIENCE was his adventure.

In a Golden Age of particle physics, Ernest Orlando Lawrence embarked on the adventure and ended up creating the model for large-scale science, winning a Nobel Prize, and infecting succeeding generations of scientists with his enthusiasm and drive.

Born a century ago on August 8, 1901, in Canton, South Dakota, Lawrence manifested his scientific spirit early. As a youngster, he didn't stop at building model gliders. Instead, he constructed an early version of a shortwave radio transmitting station. This shortwave experience was a harbinger of his future work when, in 1931, Lawrence became the first person to accelerate particles to high energies using shortwave radio techniques.

Lawrence entered college at the University of South Dakota as a premedical student. Under the tutelage of Dean Lewis E. Akeley, he switched to the study of physical science. He received his A.B. in chemistry in 1922. Years later, Dean Akeley's picture occupied a place of honor in Lawrence's office, beside portraits of scientific greats Arthur Compton, Niels Bohr, and Ernest Rutherford.

In 1928, adventure beckoned yet again, in the form of an associate professorship offered by the University of California at Berkeley. At the time of the Berkeley offer, Lawrence was already an assistant professor at Yale, where he had earned his doctorate in physics three years previously. Fellow physicist and Nobel Laureate Luis Alvarez notes in his memoirs of Lawrence, "It is difficult for one starting on a scientific career today to appreciate the courage it took for him to leave the security of a rich and distinguished university and move into what was, by contrast, a small and only recently awakened physics department." Despite the dire predictions of East Coast friends that the move would dim what otherwise appeared to be a brilliant future, Lawrence went West, hoping to have more elbow room to experiment.

By now, the story of how Lawrence chanced upon the sketch that would lead him to invent the cyclotron has become a part of scientific lore, akin to the stories of Newton and his falling apple and Archimedes and his overflowing bathtub. One evening in February 1929, while browsing through periodicals in the library, Lawrence came upon an obscure German publication with an article by physicist Rolf Wideroe, detailing a theory for ion acceleration. He didn't actually read the article—"It was in German, and I didn't read German

"E. O. Lawrence was a pathfinder not just for Lawrence Livermore and Berkeley laboratories. He created the model for large-scale science throughout the world."

—Bruce Tarter

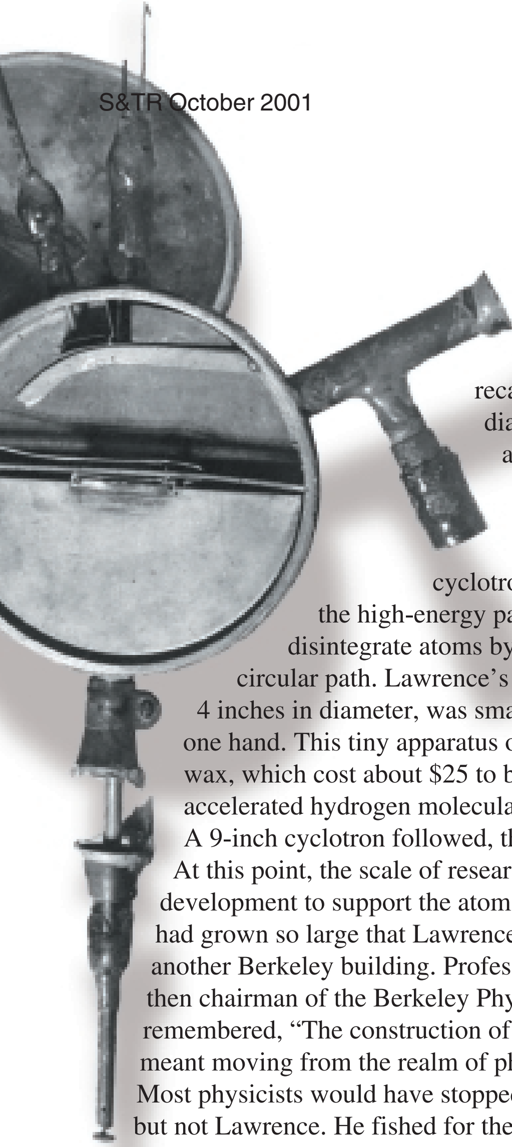

well," Lawrence recalled—but one of the diagrams drew his attention. From that single diagram, Lawrence sketched out the core design of a cyclotron—a way of producing the high-energy particles needed to disintegrate atoms by "pushing" them in a circular path. Lawrence's first cyclotron, all of 4 inches in diameter, was small enough to hold in one hand. This tiny apparatus of brass and sealing wax, which cost about $25 to build, successfully accelerated hydrogen molecular ions to 80,000 volts. A 9-inch cyclotron followed, then an 11-inch one.

At this point, the scale of research and the engineering development to support the atom-smashing projects had grown so large that Lawrence had to be moved into another Berkeley building. Professor Raymond Birge, then chairman of the Berkeley Physics Department, remembered, "The construction of the larger [cyclotrons] meant moving from the realm of physics into engineering. Most physicists would have stopped with what they know, but not Lawrence. He fished for the big ones."

In his Radiation Laboratory (known as the Rad Lab), Lawrence integrated both theoretical scientists and engineers into his projects, and the laboratory became the prototype of the big laboratories that would follow. By 1939, the Rad Lab featured a 37-inch cyclotron that, in addition to being used for exploring nuclear physics, was also being used in a radical new treatment for cancer. Lawrence was quick to see the possibilities of the cyclotron beyond pure physics. He worked alongside medical doctors, chemists, biologists, and engineers to create uses for the product radioisotopes.

On November 10, 1939, Lawrence was awarded the Nobel Prize in physics. At the time of the announcement, World War II had just broken out in Europe, and Lawrence began focusing his efforts—as so many other prominent scientists did at the time—on the war effort. Lawrence helped establish the radar program at the Massachusetts Institute of Technology and the sonar development program for antisubmarine warfare in San Diego in 1941. Back at the Radiation Laboratory, he converted the 37-inch cyclotron into a mass spectrometer, using it to successfully separate small amounts of uranium-235 from natural uranium. Lawrence's work—along with that of J. Robert Oppenheimer—greatly contributed to creating the atomic bomb. In early 1945, mere months before the first bomb was dropped, Lawrence said, "The atomic bombs will surely shorten the war, and let us hope that they will effectively end war as a possibility in human affairs."

After the war, Lawrence focused on basic scientific research at the Rad Lab until 1949, when the Soviet Union exploded its first nuclear device. Lawrence again turned his attention to national security issues. At the time, there was already a debate within government, military, and scientific

"Scientific achievement is rooted in the past, is cultivated by contemporaries, and flourishes only in a favorable environment. No individual is alone responsible for a single stepping stone along the path of progress."

—Ernest O. Lawrence

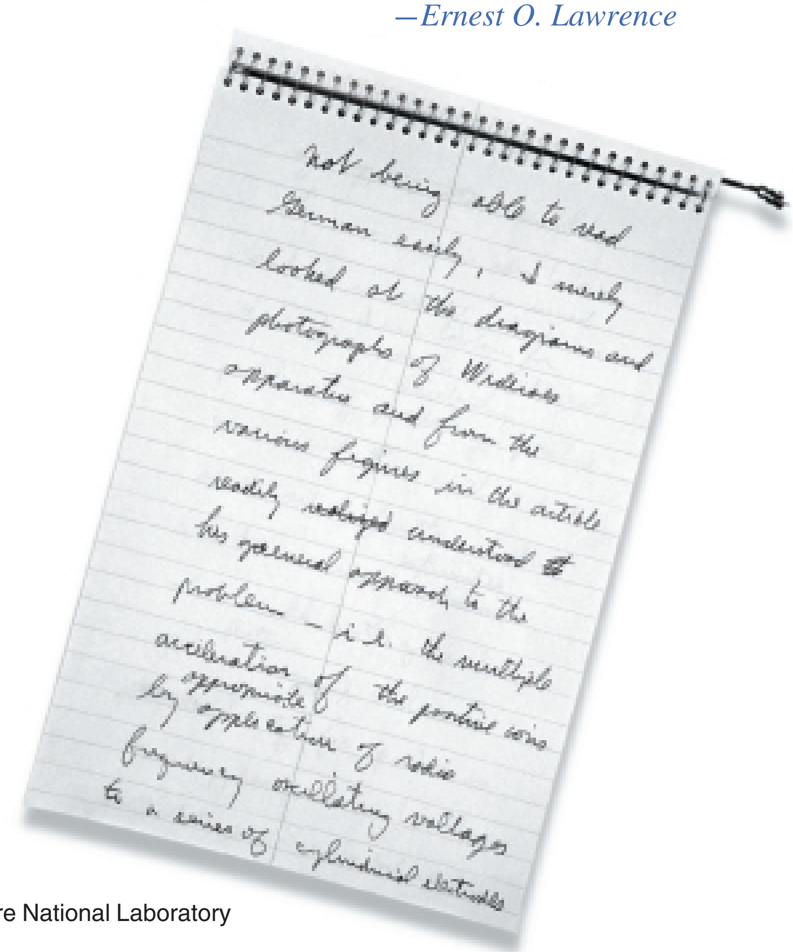

circles over the need for a second weapons laboratory. Physicist Edward Teller, who had made key contributions to the development of thermonuclear weapons while at Los Alamos, was foremost among those pressing for such a facility. Teller's rationale was twofold: the Los Alamos thermonuclear program was not proceeding at a fast enough pace, and a second laboratory would provide a source of mutual competition.

In 1951, Thomas Murray, a member of the Atomic Energy Commission (or AEC, the forerunner to the present Department of Energy) contacted Lawrence to discuss the issue of a second laboratory. Lawrence responded favorably and asked Herbert York, who at the time was one of his postdoctoral students, to survey the science community on the need for a second nuclear weapons research laboratory. Convinced that such a laboratory was needed, Lawrence then urged the AEC and the Joint Congressional Committee on Atomic Energy to set up a second weapons laboratory. He offered the Rad Lab's satellite site in the Livermore Valley as a possible location as well as his personal oversight of this new project. Livermore Laboratory was established in 1952 as a branch of the University of California Radiation Laboratory. Herbert York was its first director, overseeing daily operations at the site. Teller also remained in Livermore, working with Livermore scientists to build the nuclear program. Lawrence split his time and effort between Berkeley and Livermore. York recalls that Lawrence "always wanted to hear what was going on. We would spend one or two hours just walking around the Laboratory, visiting every place. The shops, the chemistry laboratories, the drafting rooms—there was no place he didn't want to go and where he wouldn't just stop and talk with anybody and ask, 'What are you doing right now?' and expect an answer."

Even with two laboratories to run, Lawrence found time for his family: his wife Molly and their six children. Robert Lawrence, one of his sons, recalled that his father was an inveterate tinkerer: "He built us all sorts of things growing up—scooters, wagons, other toys. … He asked us one day what he should build next, and we all said 'Build a color TV!' And so he would play with the tube on weekends, and he came up with a better, cheaper TV tube than the ones on the market. And he did it for fun, on the weekend."

Lawrence died on August 27, 1958. By the time of his death, at age 57, he had won virtually every major award in his field, ushered in a new age of physics, and developed a new way of doing science. Soon after his death, the University of California regents renamed the Berkeley and Livermore laboratories to honor him. In discussing Lawrence and the evolution of his laboratories, Herb York concluded, "I think if Lawrence were to visit the Lab today, he'd take the same 'gee whiz' attitude that he took 50 years ago. His lab has evolved in a perfectly natural way—the scope is wider, but the science is still an adventure."

—Ann Parker

As physicist Luis Alvarez notes, "One of the greatest difficulties one encounters in writing of Ernest Lawrence's career is that so much must be omitted in order to keep the account within reasonable bounds."

For more information about Ernest O. Lawrence, his life and times, see the following:
www.llnl.gov/llnl/history/eolawrence.html
www.lbl.gov/Publications/Currents/Archive/Aug-10-2001-EOL/TheMan.html
www.nobel.se/physics/laureates/1939/lawrence-bio.html
www.lbl.gov/Science-Articles/Research-Review/Magazine/1981/

"Life to him seemed to be one thrill after another, but he was also always persistent and insistent!"

—Gunda Lawrence

Each month in this space we report on the patents issued to and/or the awards received by Laboratory employees. Our goal is to showcase the distinguished scientific and technical achievements of our employees as well as to indicate the scale and scope of the work done at the Laboratory.

Patents

Process for Fabricating Composite Material Having High Thermal Conductivity

Nicholas J. Colella, Howard L. Davidson, John A. Kerns, Daniel M. Makowiecki
U.S. Patent 6,264,882 B1
July 24, 2001
A process for fabricating a composite material with high thermal conductivity for specific applications, for example, as a heat sink or heat spreader for high-density integrated circuits. The composite material produced by this process has a thermal conductivity between that of diamond and copper and basically consists of coated diamond particles dispersed in a high-conductivity metal, such as copper. The composite material can be fabricated in small or relatively large sizes using inexpensive materials. The process consists, for example, of sputter-coating diamond powder with several elements—including a carbide-forming element and a brazeable material—compacting them into a porous body, and infiltrating the porous body with a suitable braze material, such as a copper–silver alloy, thereby producing a dense diamond–copper composite material with a thermal conductivity comparable to synthetic diamond films at a fraction of the cost.

High Voltage Photovoltaic Power Converter

Ronald E. Haigh, Steve Wojtczuk, Gerard F. Jacobson, Karla G. Hagans
U.S. Patent 6,265,653 B1
July 24, 2001
An array of independently connected photovoltaic cells on a semi-insulating substrate contains reflective coatings between the cells to enhance efficiency. A uniform, flat-top laser beam profile is illuminated upon the array to produce electrical current of high voltage. An essentially wireless system includes a laser energy source being fed through optical fiber and cast upon the photovoltaic cell array to prevent stray electrical signals before the current from the array is used. Direct bandgap, single-crystal semiconductor materials, such as gallium arsenide, are commonly used in the array. The system is useful where high voltages are provided into confined spaces such as explosive detonation areas, accelerators, photocathodes, and medical appliances.

Monolithic Laser Diode Array with One Metallized Sidewall

Barry L. Freitas, Jay A. Skidmore, John P. Wooldridge, Mark A. Emanuel, Stephen A. Payne
U.S. Patent 6,266,353 B1
July 24, 2001
A monolithic, electrically insulating substrate that contains a series of notched grooves is fabricated. The substrate is metallized so that only the top surface and one wall adjacent to the notch are metallized. Within the grooves is a laser bar, an electrically conductive ribbon or contact bar, and an elastomer that secures or registers the laser bar and ribbon (or contact bar) firmly along the wall of the groove that is adjacent to the notch. The invention includes several embodiments for providing electrical contact to the corresponding top surface of the adjacent wall. In one embodiment, after the bar is located in the proper position, the electrically conductive ribbon is bent so that it makes electrical contact with the adjoining metallized top side of a heat sink.

Discrimination of Porosity and Fluid Saturation Using Seismic Velocity Analysis

James G. Berryman
U.S. Patent 6,269,311 B1
July 31, 2001
The method of the invention is employed for determining the state of saturation in a subterranean formation using only seismic velocity measurements (for example, shear and compressional wave velocity data). Seismic velocity data collected from a region of the formation of like solid material properties can provide relatively accurate partial saturation data derived from a well-defined triangle plotted in a $(\rho/\mu, \lambda/\mu)$ plane. When the seismic velocity data are collected over a large region of a formation having both like and unlike materials, the method first distinguishes the like materials by initially plotting the seismic velocity data in a $(\rho/\lambda, \mu/\lambda)$ plane to determine regions of the formation having like solid material properties and porosity.

Awards

David K. Johnson, leader of the Site Utilities Division at the Laboratory, was elected to the **Electrical Engineering Technical Advisory Committee** of the California Board of Registration for Professional Engineers and Land Surveyors. He is one of two electrical engineers selected to serve.

The board of registration is part of the Department of Consumer Affairs and regulates professional engineering and professional land surveying practice in the state. "One of my jobs is to investigate technical issues for the board and advise them on electrical engineering issues," Johnson says. "One thing the committee has talked about is the electricity crisis in the state, and I imagine that there will be technical issues that might be related to that."

Johnson received a B.S. in electrical engineering from Iowa State University and an M.B.A. in marketing from the University of Iowa. He is also the recipient of the Vice President's National Performance Award (the Hammer Award) for saving money in government.

Each year, the **American Welding Society** honors individuals who have made valuable contributions to the field of welding research, to the advancement of the welding industry, and to the society as a whole. At the society's 2001 awards luncheon held recently in Cleveland, Ohio, Laboratory employees **John Elmer** and **Joe Wong** received the **William Spraragen Memorial Award** for coauthoring, with two others, the best paper published in the Research Supplement of the *Welding Journal*, "Evolution of Titanium Arc Weldment Macro and Microstructures—Modeling and Real Time Mapping of Phases."

Elmer is a group leader for Livermore's Material Joining Group, which is responsible for electron and laser beam welding, vacuum brazing, and diffusion bonding. He also is a principal investigator for a DOE Office of Basic Energy Sciences research program to investigate phase transformation kinetics during synchrotron radiation x-ray diffraction measurements of arc welds. Elmer has a Ph.D. in metallurgy from the Massachusetts Institute of Technology. He was named an American Welding Society Fellow in 2000.

Wong has B.Sc. degrees in pure and applied chemistry and physical chemistry from the University of Tasmania, Australia, a Ph.D. in physical chemistry from Purdue University, and a D.Sc. from the University of Tasmania. He joined the Laboratory in 1986 and is currently a senior staff chemist. Since 1995, he has also been consulting professor at the Stanford Synchroton Radiation Laboratory. He has published over 175 journal articles, coauthored a textbook, and been granted seven U.S. patents.

Abstract

Sharing the Power of Supercomputers

Lawrence Livermore created the Multiprogrammatic and Institutional Computing (M&IC) Initiative in 1996 to ensure that all Livermore programs and researchers have the possibility of accessing high-performance computers. The initiative, a partnership between the Laboratory and its research programs, serves more than a thousand users, including outside collaborators, with some of the most powerful computers available. Under the M&IC Initiative, the Livermore Computing Center has acquired increasingly more powerful clusters, or groups, of computers, including the TeraCluster2000 and the unclassified ASCI Blue and ASCI Frost supercomputers. These machines make it possible for researchers to perform close to teraops-scale (1 trillion mathematical operations per second) parallel computing to achieve breakthrough scientific discoveries. Parallel computing attacks huge mathematical problems with a number of identical processors simultaneously sharing a computational task. The resources provided by the M&IC Initiative are permitting researchers to generate simulations that, in many cases, were never before attempted for lack of computing power.

Contact:

Mike McCoy (925) 422-4021 (mccoy2@llnl.gov) or

Greg Tomaschke (925) 423-0561 (tomaschke1@llnl.gov).

Welding Science Revealed in Real Time

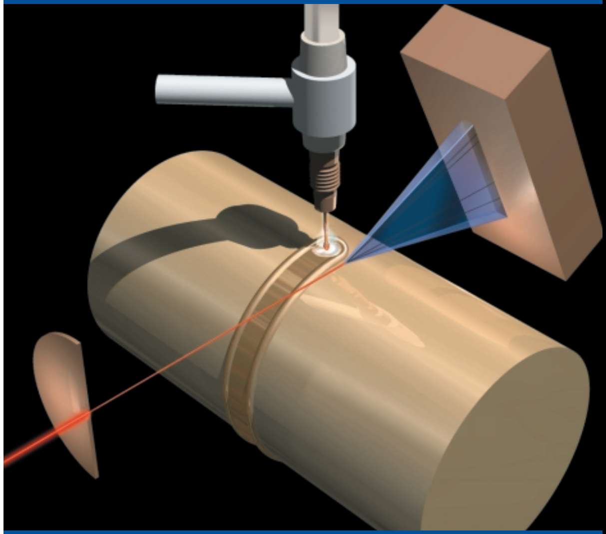

Livermore material scientists are studying the science of the welding process as it occurs.

Also in November

- ***Crosswell electromagnetic imaging is being used to monitor the effectiveness of injecting carbon dioxide underground to enhance oil recovery.***

- ***New fusion fuel targets are being designed to exploit two promising techniques—heavy-ion fusion and fast ignition—for producing commercial fusion power.***

- ***The atomic structure of water molecules at the surface of a water jet is being probed with soft x rays.***

Coming Next Month

U.S. Government Printing Office: 2001/683-051-70003

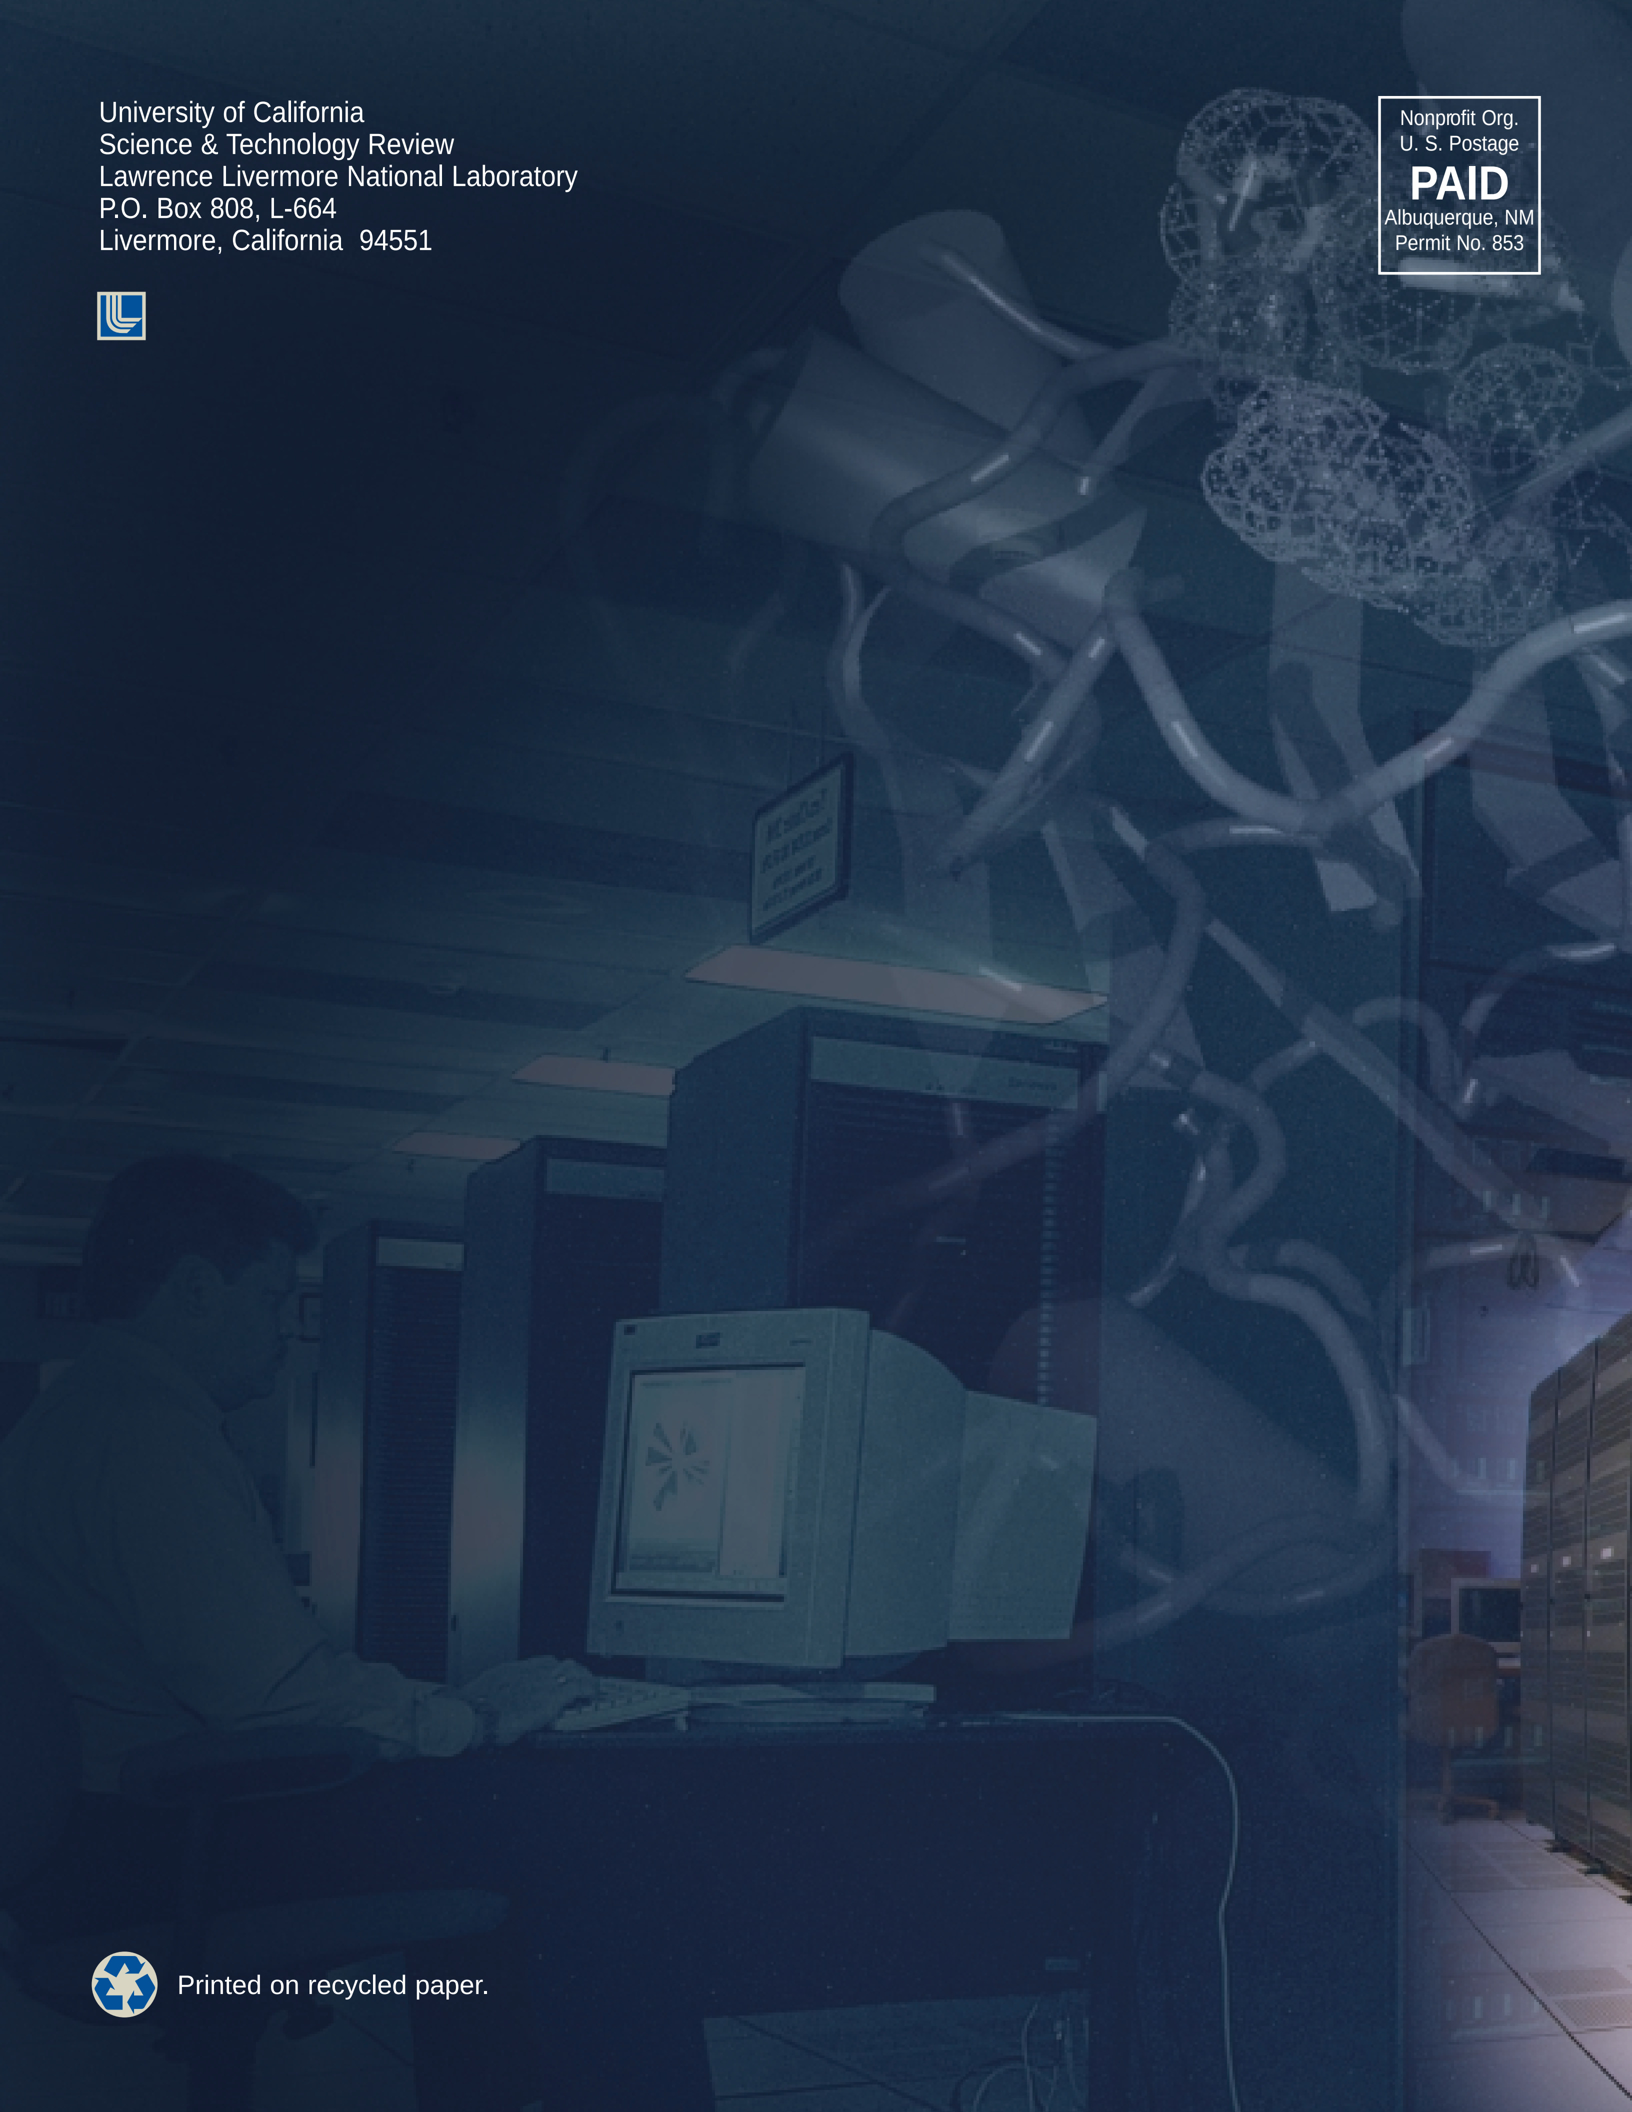

University of California
Science & Technology Review
Lawrence Livermore National Laboratory
P.O. Box 808, L-664
Livermore, California 94551